IBM® BladeCenter™ Deployment Guide
Part 3 - Blades

Chris Gillen, IBM

Sam Litenatsky, IBM

Kevin Conover, IBM

Table of Contents

Table of Tables

Table of Figures

1 Introduction

1.1 Overview

The IBM BladeCenter Deployment Guide is intended to present the general steps necessary to successfully deploy IBM blades in the IBM® eserver BladeCenter™. Its primary focus is to detail the steps required for setting up remote boot on IBM Blades. The information provided in this document is based on technical observations.

Due to the breadth of content related to this topic, this guide has now been split into 3 parts.

1. Management Module – information specific to the initial setup and ongoing usage of the Management Module.
2. Embedded Switches – information related to the configuration of the embedded Ethernet and Fibre channel switches.
3. Blades – blade configuration and installation information with a primary focus of installing Linux in a boot to SAN environment.

Remote boot, root boot, or boot from SAN, is the name referred to the server configuration where the server operating system is installed on a logical drive (LUN) that does not reside inside the server chassis. This document will describe in detail the process that one must go through to setup a remote boot for HS20 or HS40 blade servers. The blade servers utilize the IBM BladeCenter Fibre Channel Expansion Card connected to the BladeCenter Fibre Channel Switch Modules via the BladeCenter backplane. The switch modules in turn are connected to the core fibre channel switches that host the SAN Fabric.

This document is a supplement the IBM® eserver BladeCenter™ publications provided with the BladeCenter products. In addition to the publications being provided with the BladeCenter products, they also are available in Portable Document Format (PDF). The latest versions of the publications are available from the IBM web site at http://www.ibm.com/pc/support/.

BladeCenter publications include the following:

- *IBM eserver BladeCenter Type 8677 Planning and Installation Guide*: This manual provides physical planning information for the IBM BladeCenter products. It describes the BladeCenter components, explains deployment and installation considerations, and provides worksheets that you can use to determine the configuration, power, weight, and cabling requirements for your BladeCenter unit. (http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-53670)
- IBM eserver *BladeCenter Type 8677 Installation and User's Guide.* This document contains general information about your BladeCenter unit, including information about features, how to configure your BladeCenter unit, and how to get help. It is provided with the BladeCenter unit and is on the IBM *BladeCenter Documentation* CD.
- *IBM eserver BladeCenter Management Module User Guide*: This guide contains information about configuring the management module and managing components installed in BladeCenter chassis. (http://www-1.ibm.com/support/docview.wss?rs=0&uid=psg1MIGR-45153)
- *IBM eserver xSeries and BladeCenter Server Management* (SG24-6495): This document covers the hardware side of the IBM systems management solution including the BladeCenter management module. It is available for download from the IBM Web site at http://www.ibm.com/redbooks/.

1.2 BladeCenter Subsystems

The various BladeCenter Subsystems include:

- Server Blades - provide the processors, memory, hard disk(s) and firmware of an industry standard server. In addition, they will normally incorporate keyboard, video, and mouse interfaces, and an onboard service processor.
- Management Modules – incorporate the chassis-level management processor and Keyboard/Video/Mouse (KVM) access for the blades.
- Switch Modules - provide multiple high-speed Ethernet ports in Switch Module slots 1 and 2 for use by all of the Server Blades to interconnect internally as well as to external network devices. Switch

Module slots 3 and 4 are used for Fibre Channel or other switch types and connect to the Server Blades as well as external devices.

- Power - Four modular power supplies provide primary and backup power to all BladeCenter subsystems. Power supplies 1 and 2 are responsible for powering blade 1 through 6 plus all of the BladeCenter components such as the switches, management modules, media tray, etc, and power supplies 3 and 4 are responsible for power blades 7 through 14.
- Cooling - Dual, modular blowers provide forced-air cooling of BladeCenter components and subsystems.
- Media Tray - Removable media subsystem that is shared by all blades. Includes CD/DVD and floppy media.

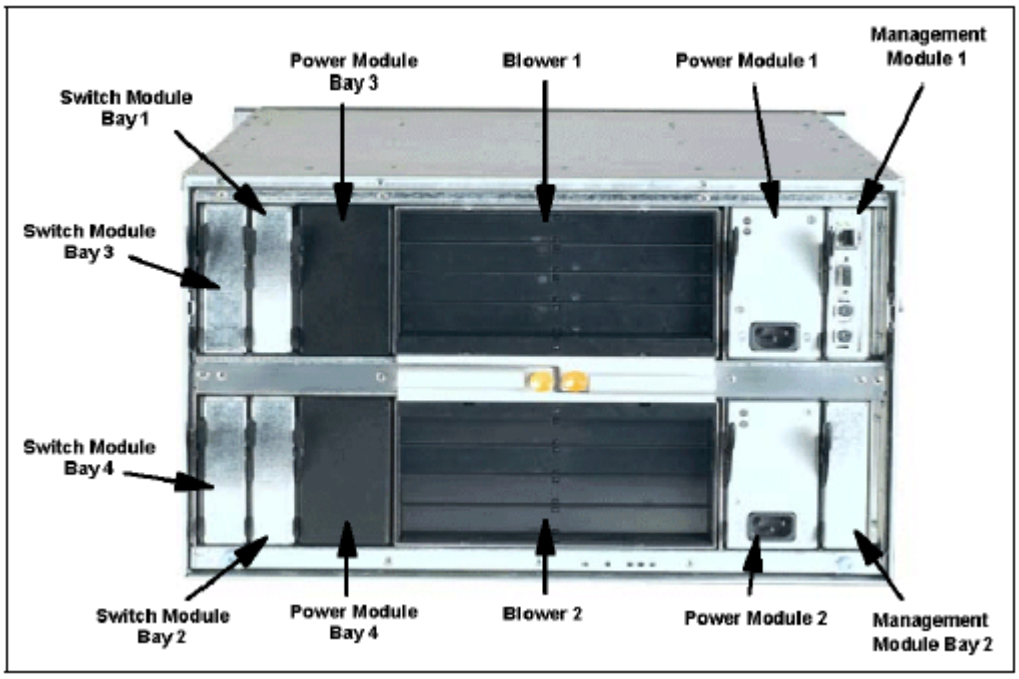

Figure 1 BladeCenter Chassis – Rear View

1.3 Notes

This is currently a living document and will continue to grow as we have more experience with IBM Blades in the boot-to-SAN environment. This document is intended as a personal productivity tool. It is not intended to be comprehensive, and is provided for guidance only, on an 'as is basis' without warranty of any kind. Please be aware that its contents have not been certified by IBM**.**

⚠⚠ If you are installing Linux, please pay close attention to the issues listed in Section 4.5.

1.3.1 Contributors

We would like to acknowledge the contributions of the following people:

1. Richard Gebhardt, IBM Lab Services – Appendix E Modify or Create the Post-install Driver Diskette Image, portions of Section 4 Installation Process for Red Hat Linux.
2. Fred Rabert, Nortel Networks – Assistance with Section on Nortel switches plus BroadCom Driver installation instructions.

Installation and Configuration

2 Preparatory Steps

The BladeCenter Chassis components should be prepared by updating the firmware and assigning the appropriate IP addresses to the management and switch modules. The IP addresses of the various modules as well as the userids and passwords for each will be needed during the installation. A sample worksheet can be found in 1.Appendix A.

The operating systems that are supported in BladeCenter boot from SAN configurations are Linux RedHat Advanced Server 2.1, 3.0 and 4.0; Microsoft Windows 2000 (W2K) Advanced Server, Microsoft Windows 2000 Server, Microsoft Windows 2003 (WS03) Enterprise Edition, Web Edition and Server.

> **Note:** Since the storage subsystem provides multiple paths to the server, multiple views of the same logical disk will be presented. Multipathing software controls these multiple views by interposing between the operating system and the hardware. Since the multipathing software is not available for the initial boot, only one path may be presented to the server for the initial boot. This process is described in detail later in this document.

2.1 Network Configuration

The two switch modules provide redundant paths to each of the server blades for reaching other server blades within the same chassis or for communication to external devices via the network infrastructure. (Figure 2 BladeCenter internal network) The Ethernet switch modules can be managed via telnet or a Web interface. Both the Web and telnet interfaces can be started by accessing the switch directly or starting a session from the management module's Web interface.

Note: The external ports on the switch must be enabled, and enabled for management as well, to manage the switch directly without going through the management module. Because the management interface is on VLAN 4095 and cannot be moved, a separate interface from the management interface must be created with an IP address valid on the production network to manage the switch directly.

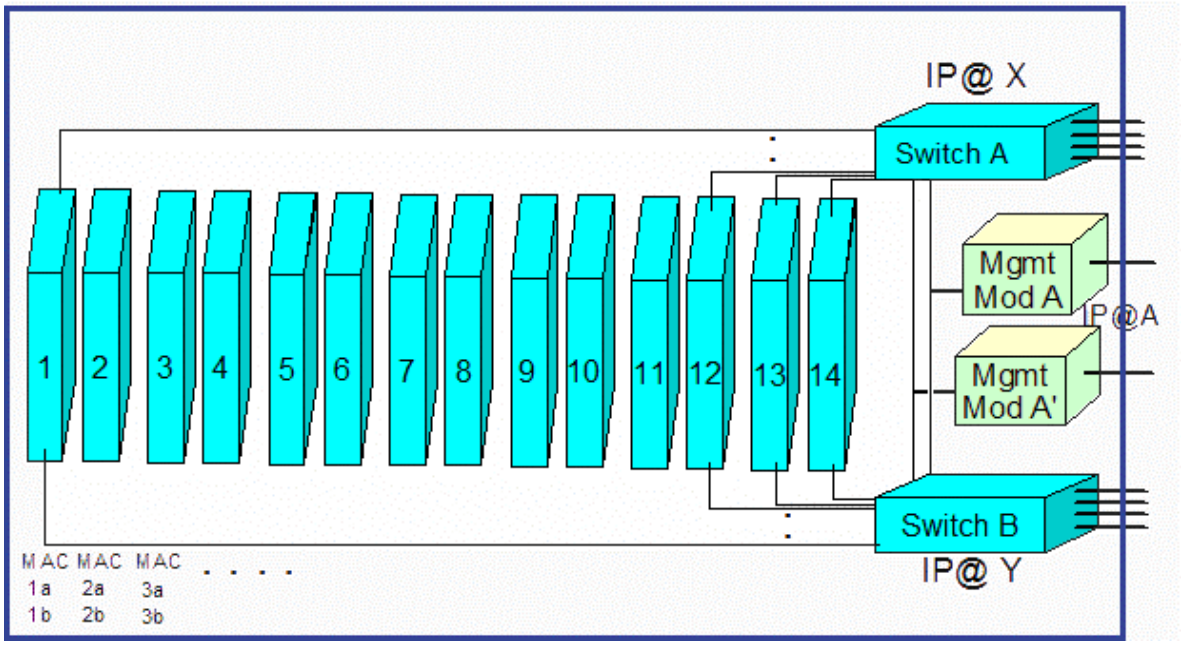

Figure 2 BladeCenter internal network

Note: For switch communication through the management module external Ethernet port, the switch module internal network interface and the management module internal and external interfaces must be on the same subnet.

2.2 BladeCenter Management Module

To use the Management Module to perform tasks in the coming sections remotely, first do the following: Open a browser and input the IP address of the MM. The default IP address of the external port of the management module is 192.168.70.125, the default subnet mask is 255.255.255.0, and the default hostname is MM*xxxxxxxxxxxx* where *xxxxxxxxxxxx* is the burned-in MAC address. It is possible to connect directly into the management module by using a crossover cable between a laptop and the Ethernet port of the management module. Change the ip address of the laptop to an address on the same subnet as the management module. 192.168.70.124 is a good choice.

- The logon screen will display, and you will be prompted to log onto the MM web server. The default user ID and password are `USERID` and `PASSW0RD` (`0` is a zero). These are case sensitive.
- Once you enter your user ID and password, you will presented with a welcome window that displays where you are connected and, for security purposes, where you set the duration that your connection can be idle before being automatically disconnected.
- Click <Continue>. The BladeCenter Management and Configuration window opens. From the Management and Configuration program main menu, select settings that you want to view or change.

2.3 Remote Control Hints and Tips

Note: All of the subsequent tasks described in the document can be performed either from direct KVM connection at the BladeCenter or via the Management Module remote console. In most cases, instructions will be provided assuming access via the Management Module.

1. Sun JVM Version 1.4.2 is the minimum required for remote control function. The Microsoft JVM is not supported. It may work some of the time, but Remote Disk tends to act flaky with the MSJVM. To check this in Internet Explorer, go to Tools-Options, and then click on the Advanced tab.

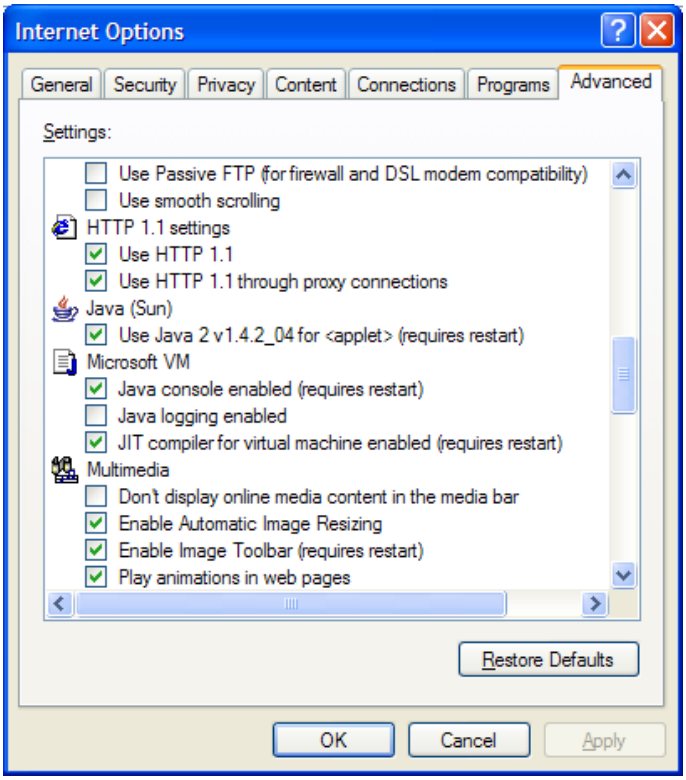

Figure 3 JVM Selection in IE

2. Repeating Keystrokes - in order to remove the repeating keystroke problem, you need to check the "Disable key repeats" box and click Save.

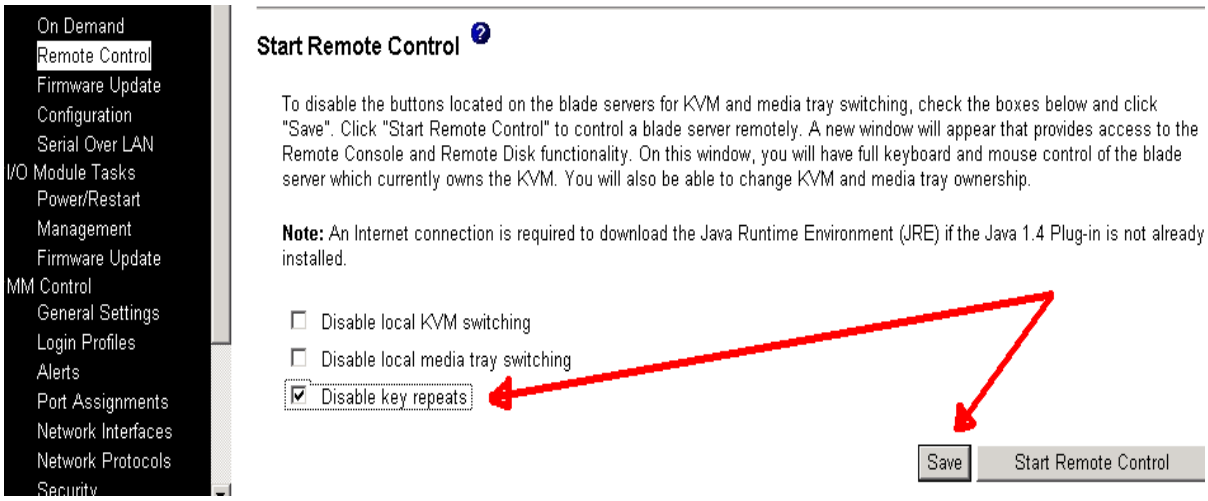

Figure 4 Disable key repeats

3. Start a Remote Control of the target blade by clicking on **Remote Control** under **Blade Tasks**. Ensure that the media tray and KVM selections are for the target blade)

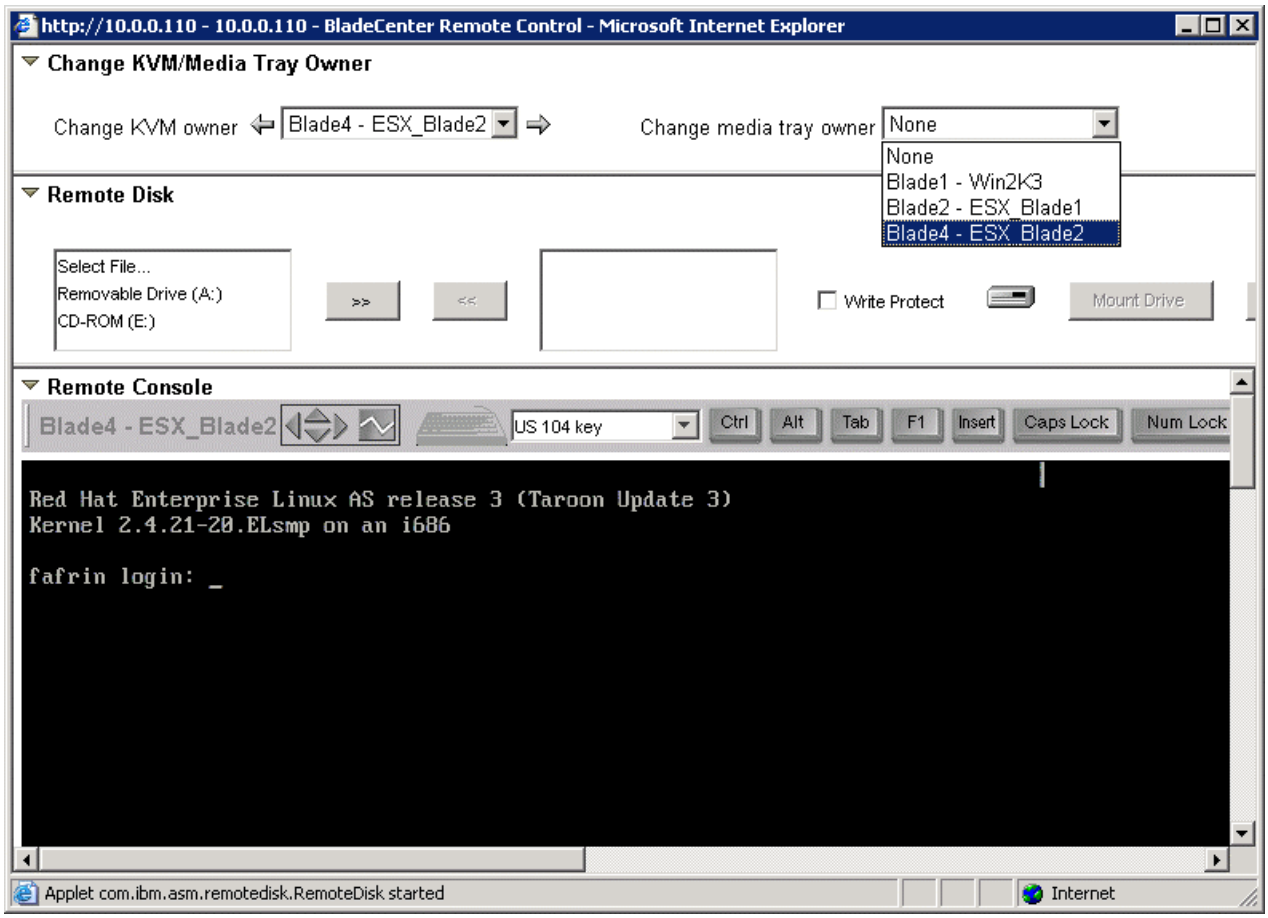

Figure 5 KVM and Media Ownership

4. The Remote Control window is a pop up window. If you have a popup blocker installed, Remote Control will be blocked, so make sure popups are allowed for the management module IP. Note: Remote Control does work with Firefox, so if you are having trouble with the popup blocker in Windows XP SP2, Firefox may be a viable alternative.
5. The Remote Control window includes a tool bar (Figure 6) with "sticky keys" for passing "`ctrl`", "`alt`", and others keystrokes to the target system. This is particularly important when trying to enter setup screens, such as the QLogic setup which is accessed by hitting "`ctrl-Q`". Click on the sticky "`ctrl`" then hit Q when you see the message during POST.
6. If there is no keyboard response in Remote Control, confirm that the Remote Control window has focus. Click on the keyboard widget in the toolbar. There is also a known issue, fixed in the 73G release of the Management Module firmware, where the "`ctrl`" or the "`alt`" keys become stuck on. Clicking the icons on the toolbar twice should clear this state and may restore keyboard response.

Figure 6 Remote Control Window

7. If, during a Linux installation, you see the message "No Video Available" (shown in Figure 6) then most likely "nofb" was not specifying as a boot option. See Section 4.3 for more information.
8. When switching from managing one chassis to another, it is recommended that you close your browser after logging off from the first chassis, especially if you were using Remote Control in the first session. Otherwise, you may encounter errors with the Remote Disk applet (middle frame, Figure 5) in the second session such as "remotedrive.dll is not accessible" "exception: java.lang.NullPointerException." "Initialization error. Library not installed" "Unsatisfied Link Error – Applet com.ibm.asm.remotedisk.RemoteDisk not inited."

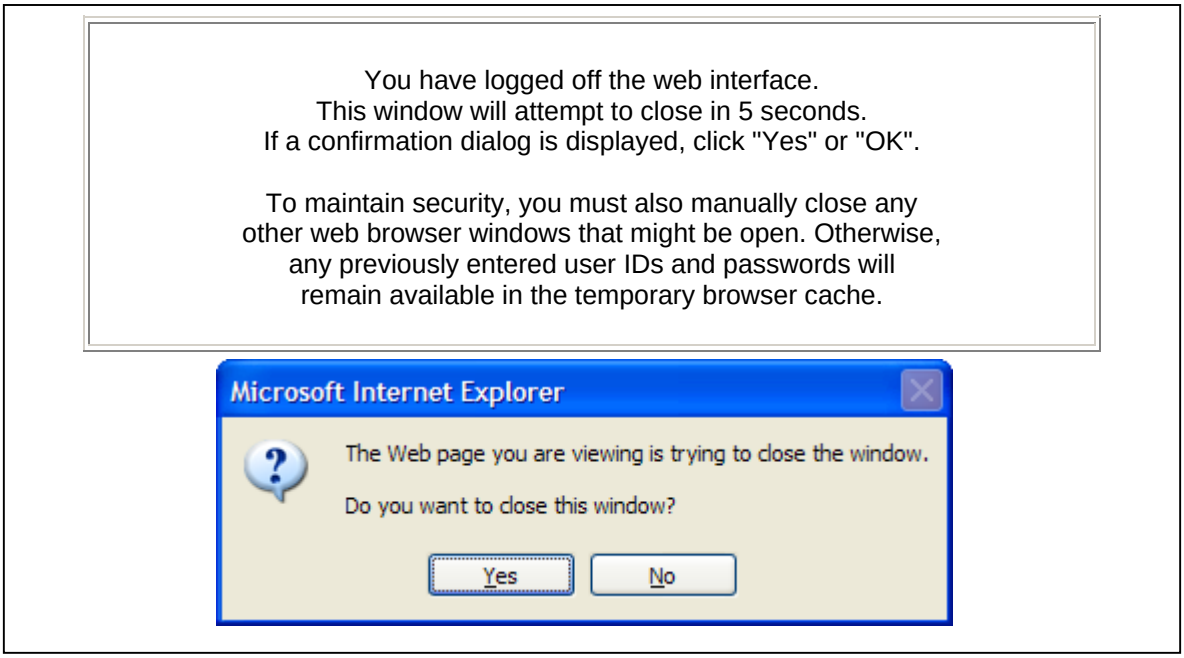

Figure 7 Logoff and Close the Browser Messages

9. Unexpected DLL version number error when starting Remote Control.,

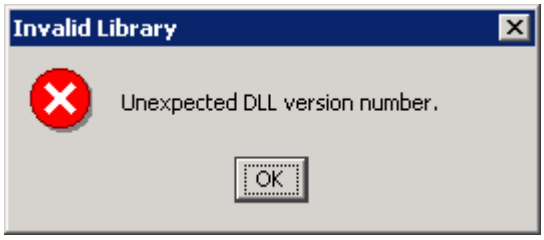

Initializing RemoteDisk v2.0
Attempting to load remotedrive.dll
Remote Disk library successfully loaded.
DLL reported: 12. Expected: 13
Unexpected DLL version number.

If you encounter this error when the remote control window is opened, exit Internet Explorer, find the copy of remotedrive.dll on your system (most likely in c:\Program Files\Internet Exploder) and delete it.

2.4 Embedded Fibre Channel Switches

2.4.1 SAN Configuration

Shown in Figure 8 is an example of a dual path SAN configuration that we will be using as a reference in this document. Two McData directors steer traffic through separate fabrics to the IBM BladeCenters, which use Qlogic HBAs (Host Bus Adapters). Separate fabrics allow full redundancy in the storage path. The IBM Blade Server boots to drives physically located on the Hitachi subsystem over the SAN. Qlogic HBAs must be specifically configured to allow remote booting. See diagram 1 for a visual representation of this architecture.

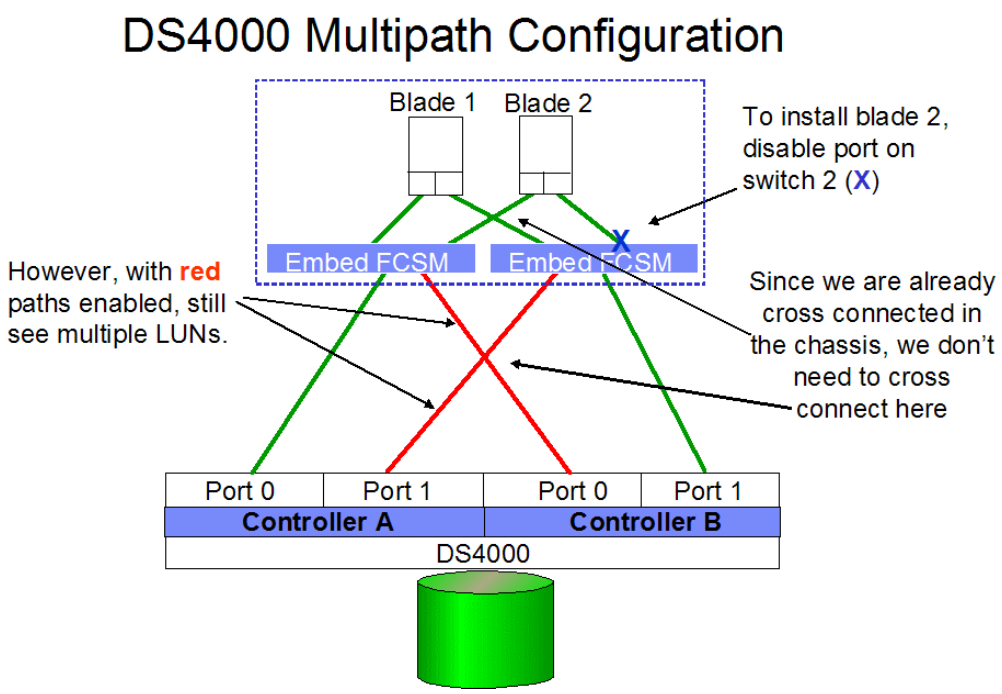

Figure 8 Storage Connection Diagram

2.4.2 Disabling ports on the QLogic switch

Here are the steps necessary to disable ports on the QLogic switch. This can be done either via a telnet session or a browser session.

2.4.2.1 Using Telnet

1. Start a telnet session with the IP address of the QLogic Switch Module.
2. After logging in, enter the command **"show port"** to view the current status of the ports.

```
        Admin       Operational  Login       Config Running     Link      Link
Port    State       State        Status      Type   Type        State     Speed
----    -----       -----------  ------      ------ -------     -----     -----
Ext1:0  Online      Online       LoggedIn    GL     F           Active    1Gb/s
Ext2:15 Online      Offline      NotLoggedIn GL     Unknown     Inactive  Auto

Bay1    Online      Online       LoggedIn    F      F           Active    2Gb/s
Bay2    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay3    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay4    Online      Online       LoggedIn    F      F           Active    2Gb/s
Bay5    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay6    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay7    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay8    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay9    Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay10   Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay11   Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay12   Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay13   Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
Bay14   Online      Offline      NotLoggedIn F      Unknown     Inactive  2Gb/s
```

Figure 9 Port Status Online

3. Enter admin mode by entering “admin start”
4. Disable the target port by entering the command “**`set port x state offline`**” where **`x`** is the number of the port. For this example, we used “**4**”
5. Confirm the status by entering “**`show port`**”again.

```
QLogic: USERID> admin start
QLogic (admin): USERID> set port 4 state offline
QLogic (admin): USERID> show port
          Admin       Operational  Login       Config Running     Link     Link
  Port    State       State        Status      Type   Type        State    Speed
  ----    -----       -----------  ------      ------ -------     -----    -----
  Ext1:0  Online      Online       LoggedIn    GL     F           Active   1Gb/s
  Ext2:15 Online      Offline      NotLoggedIn GL     Unknown     Inactive Auto

  Bay1    Online      Online       LoggedIn    F      F           Active   2Gb/s
  Bay2    Online      Offline      NotLoggedIn F      Unknown     Inactive 2Gb/s
  Bay3    Online      Offline      NotLoggedIn F      Unknown     Inactive 2Gb/s
  Bay4    Offline     Offline      NotLoggedIn F      Unknown     Inactive 2Gb/s
  Bay5    Online      Offline      NotLoggedIn F      Unknown     Inactive 2Gb/s
  Bay6    Online      Offline      NotLoggedIn F      Unknown     Inactive 2Gb/s
```

Figure 10 Port Status Offline

2.4.2.2 Using SANSurfer GUI

1. Start the SANSurfer GUI. If the fabric has not already been discovered, click Fabric-Add Fabric and enter the IP address of the target QLogic switch in the window that will pop up - That should open SAN Browser.
2. Click the handle in the left pane to drop down the switch - highlight it and the faceplate will appear in the right pane.

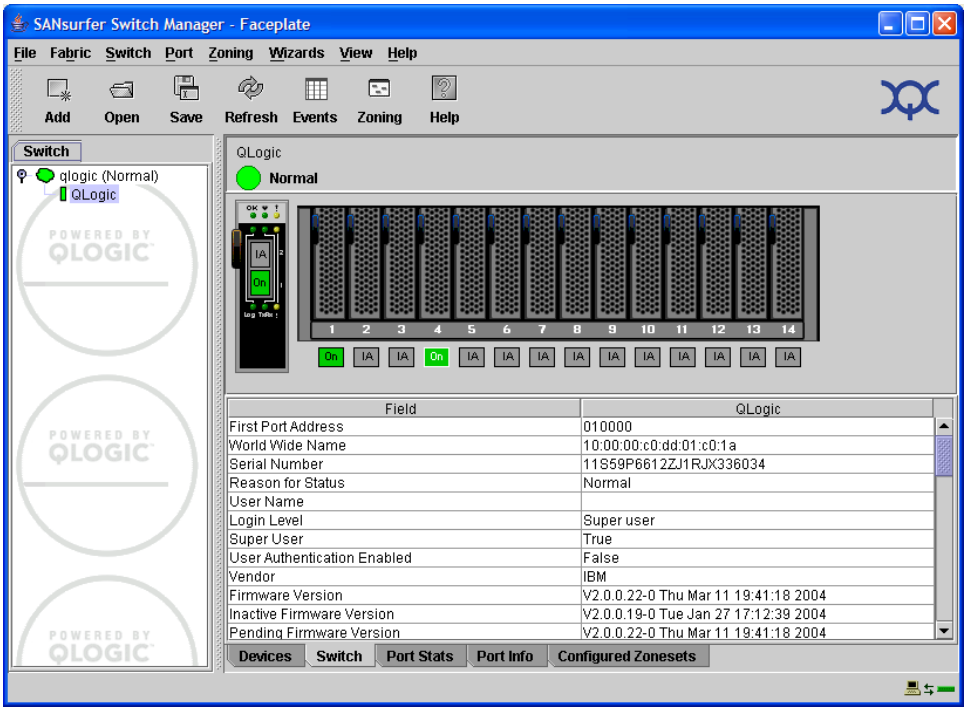

Figure 11 SANSurfer Switch Faceplate

3. Right click on the desired port and choose "port properties" from the menu.
4. At the top of port properties put a bullet next to "offline" or "down".
5. Repeat steps 1-4 and reverse step 5 to online the port.

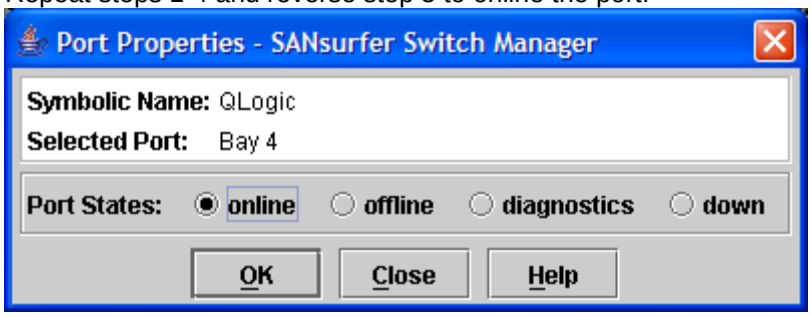

Figure 12 Port Properties

http://knowledge.storage.ibm.com/servers/storage/support/hbasearch/interop/hbaSearch.do

Parameters	Recommended settings
BIOS	Disabled
Frame size	2048
Loop reset delay	5 (minimum)
Adapter hard loop ID	Disabled
Hard loop ID	0
Execution throttle	100
Fast command posting	Enabled
>4 GB addressing	Disabled (for 32-bit systems)
LUNs per target	0
Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes Note: Enable LIP reset, Enable LIP full login, and Enable target reset parameters control the behavior of the adapter when Windows tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.
Login retry count	20 (minimum)
Port down retry count	20 (minimum)
Driver load RISC code	Enabled
Enable database updates	No
Disable database load	No
IOCB allocation	256
Extended error logging	Disabled

Table 1 Recommended settings for the QLogic QLA23xx adapter card for Windows 2000 or 2003

Blade Configuration and Installation

3 Configuring the Blades for Remote Boot

The next few steps require BIOS configuration. To enter the BIOS, follow these steps:

1. Logon to the Management Module and select **Power/Restart** under **Blade Tasks.**
2. Select the blade to be restarted, then at the bottom of the page click **Restart Blade** and "Ok" when prompted.
3. Start a Remote Control of the target blade by clicking on **Remote Control** under **Blade Tasks**. Ensure that the media tray and KVM selections are for the target blade)
4. Enter the Blade BIOS by pressing **F1** or clicking the **F1** button at the appropriate time during POST.

3.1.1 Disabling IDE devices

Since the BladeCenter servers will boot to drives located on the SAN, the locally attached drives will not be used. To boot to the Qlogic HBAs, the local IDE controllers for the IBM BladeCenter should be disabled. To disable the IDE devices in the Blade Server BIOS, use the following steps from within the BIOS setup program:

1. Select **Devices and IO ports** from the main menu
2. Select IDE configuration menu. Select **Disable** to disable the Primary and Secondary IDE controllers.

Note: in some instances, you may see the following error:

I9990303 No Boot device found

If you encounter this error, then spin up delay should be enabled on the HBA. See Section 3.2 Configuring the Qlogic HBA for details.

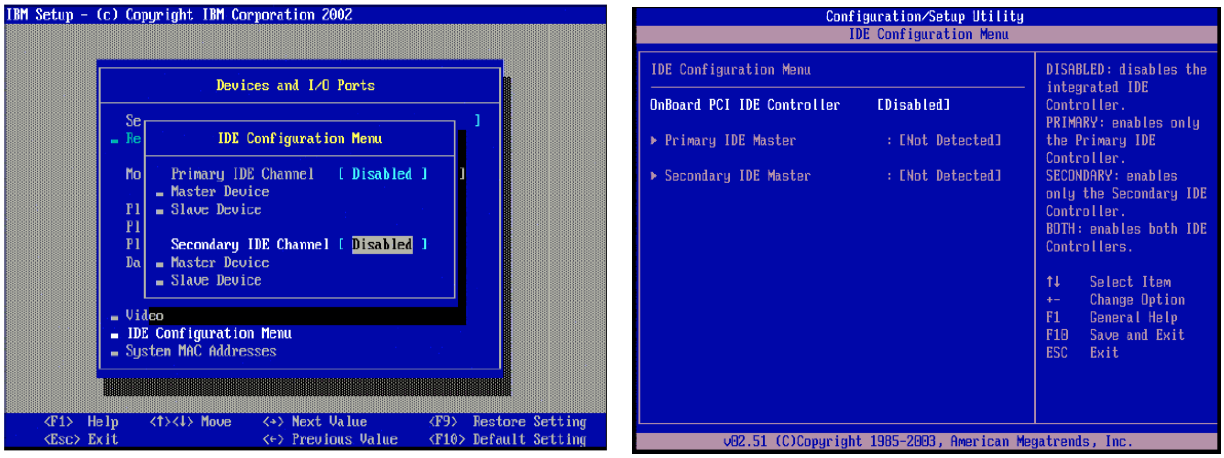

Figure 13 Disable the IDE Devices, HS20 and HS40

3.1.2 Check the Boot Order

- From main BIOS screen, access Start Options, then Startup Sequence to access the boot order and insure the devices are listed in the following order:

Device	Startup Order
CD-Rom	First Startup Device
Diskette Drive 0	Second Startup Device
Hard Disk 0	Third Startup Device
Network	Fourth Startup Device

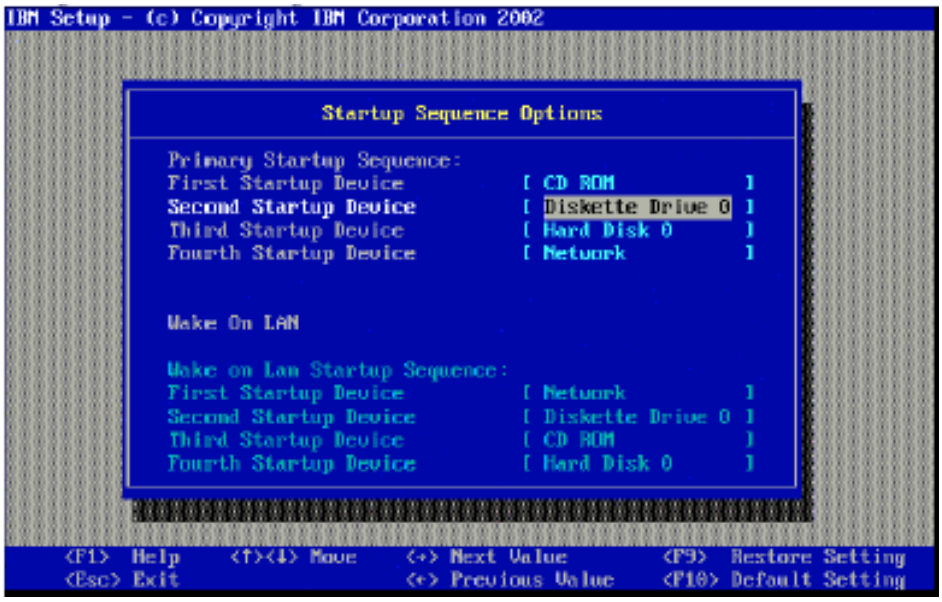

Figure 14 HS20 Boot Sequence

3. Adjusting the boot sequence for HS40 blades is slightly different. Use the arrow keys to scroll to each startup device, then hit enter and select the appropriate value in the popup window as shown in Figure 15 HS40 Boot Sequence.

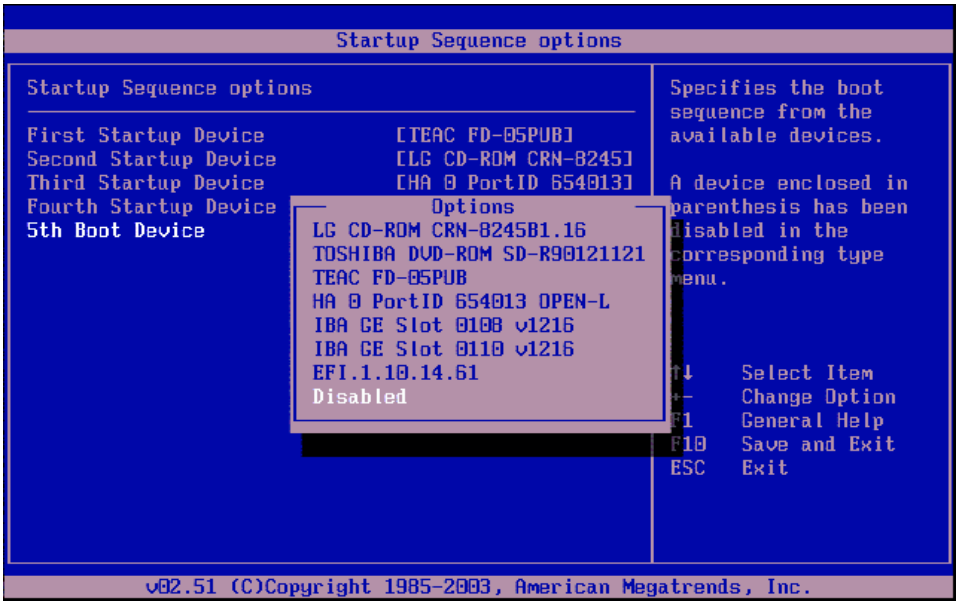

Figure 15 HS40 Boot Sequence

4. Press **ESC** to exit out of the BIOS and select the option to save changes.

> **Note:** For of HS40 blades, make sure that EFI is at the end of the boot order if present.

5. If you have a lot of blades to update in a single blade chassis, the boot order changes may be made through the Management Module and applied to any or of all of blades in the chassis as a group. (See Figure 16 Changing boot order from Management Module) Since EFI appears as an unknown device for HS40s, it is recommended that the HS40 boot sequence be modified in BIOS program.

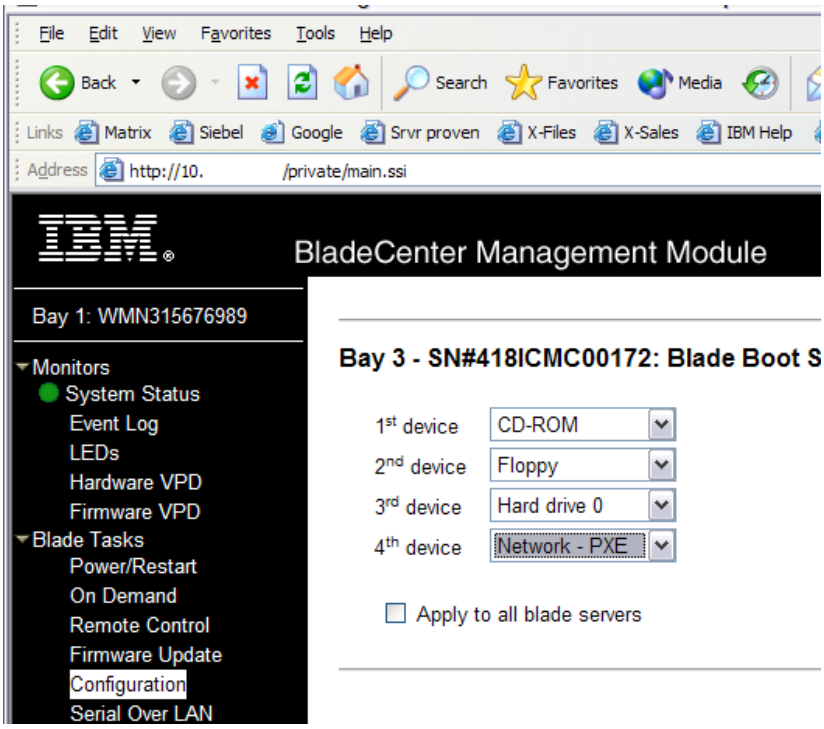

Figure 16 Changing boot order from Management Module for HS20

3.2 Configuring the Qlogic HBA

3.2.1 Acquiring HBA WWPNs (World Wide Port Names)

Each HBA has a unique WWPN by which it is identified in a SAN network. Accurate tracking of WWPNs ensures that each server has the proper storage presented to the appropriate location. The recording of the WWPN information is only required **once** for ESM to perform zoning for the SAN. If this has already been performed, skip to Section 3.2.2.

To acquire these WWPNs:

- Restart the target blade server being configured for remote boot and enter the QLogic BIOS utility using **CTRL+Q**. This must be done at a specific time during the POST. Note: you must click the "sticky" Ctrl icon on the Remote Control page when using the Management Module for this task, and then you'll only need to press "Q" on the keyboard.
- One at a time, select each HBA listed and record the WWPNs

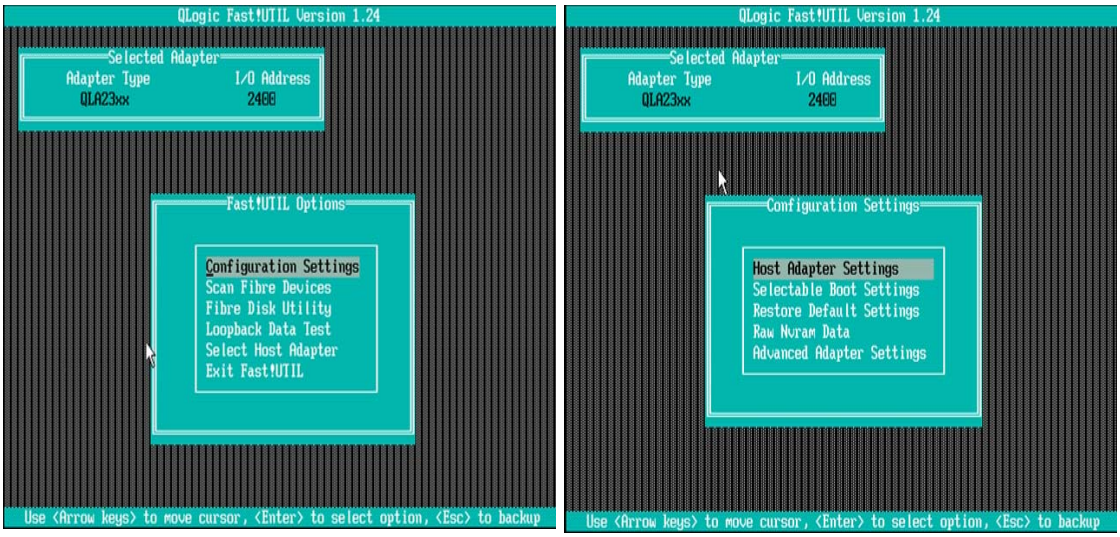

Figure 17 Configuring the QLogic HBA

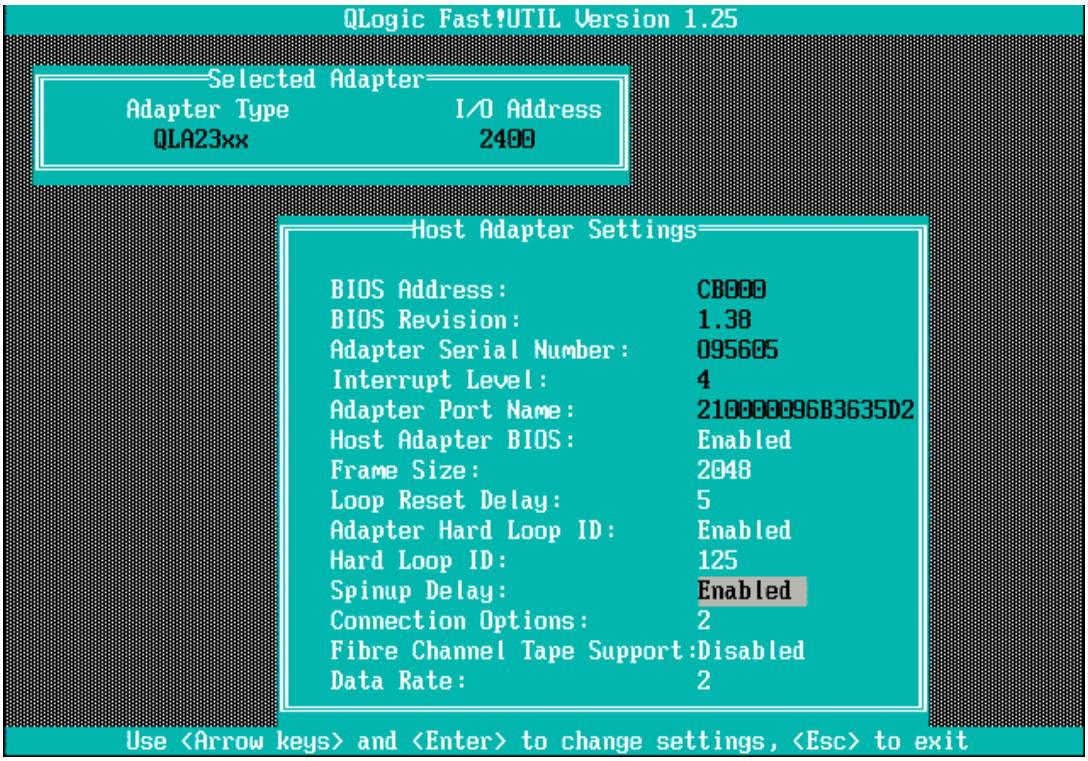

Figure 18 Host Adapter Settings

3.2.2 Configuring the path to the boot LUN

The HBA that the server will be booting from will need to be defined in the bios. At this point the HBA's WWID has been zoned on the backend storage as the owner of the boot LUN.

- If you are not already in the QLogic BIOS utility, restart the target blade server being configured for remote boot and enter the QLogic BIOS utility using **CTRL+Q** as described in Section 3.2.1.
- Select **Configuration Settings,** then **Selectable Boot Settings,** then set Selectable Boot to **Enabled**.
- Select primary boot path, press Enter, and the LUN ID should be visible as shown in Figure 20.

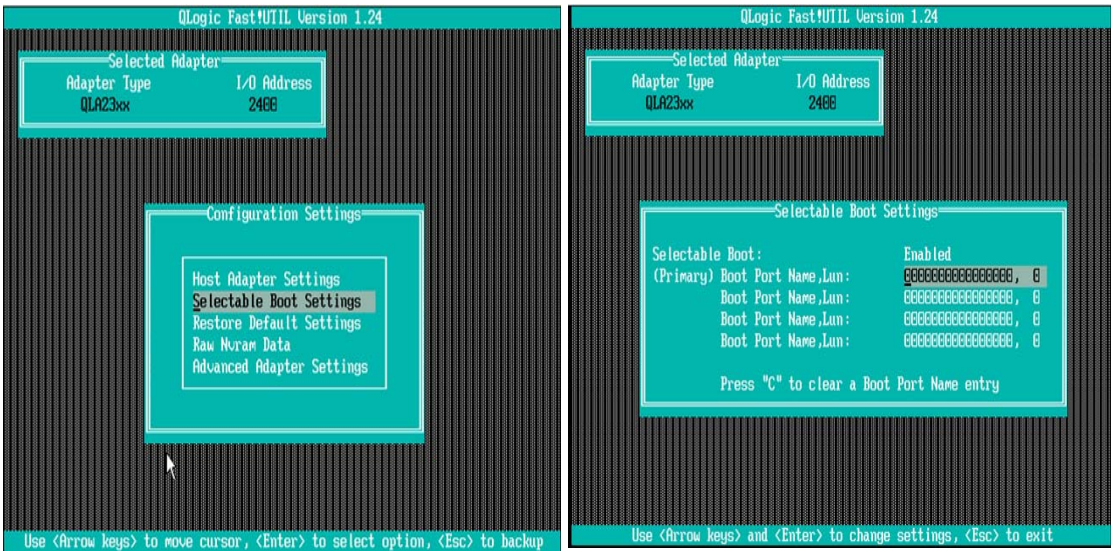

Figure 19 Enable Selectable Boot Illustration

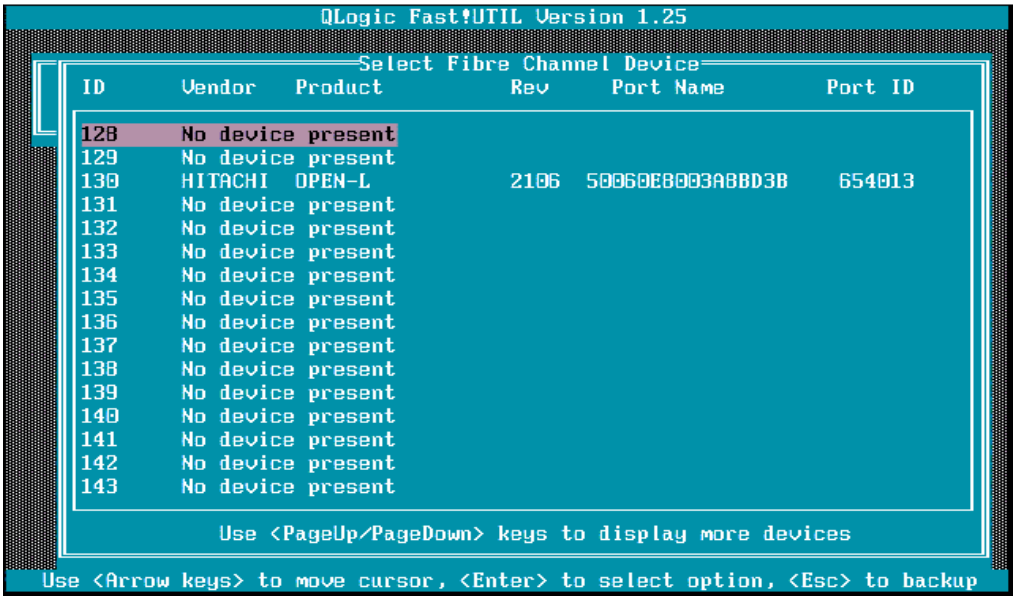

Figure 20 Bootable LUN address

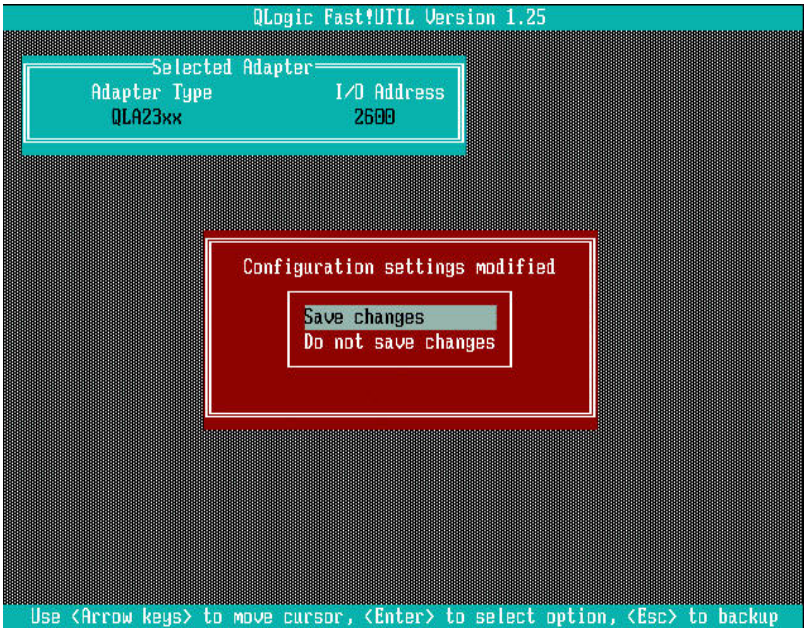

Figure 21 Save your changes

Critical Note: At this point, only one boot path is configured. During the INITIAL Operating System build, only one boot path should be active. After the initial OS build a secondary path can be added if this is supported by the storage vendor.

3.2.3 Testing the path to the boot LUN

The path to the LUN that the server will be booting from can be tested at this point.

- Select verify disk media
- Once the boot media is verified, save changes and exit.

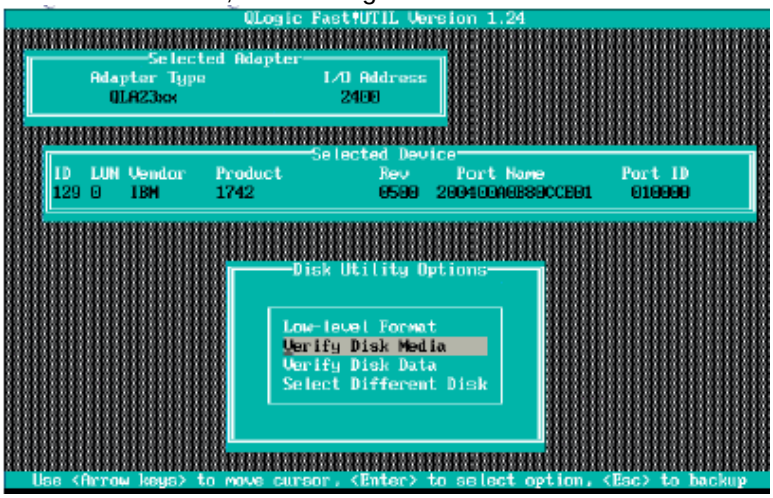

Figure 22 Confirm the media is accessible

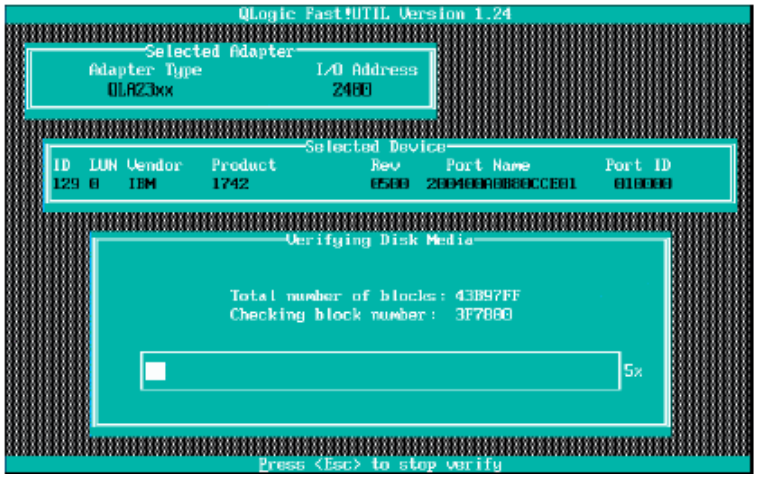

Figure 23 Verify Disk Media

Once the HBA configuration steps are completed, during the POST sequence you will notice the following text which indicates that the blade is now booting from the LUN 0 associated to the Fibre Channel Daughter Card.

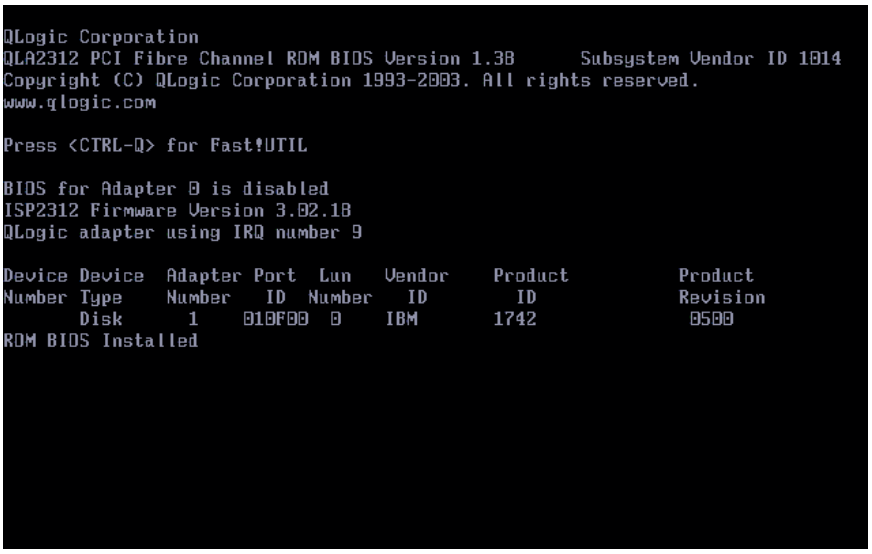

Figure 24 Blade POST with HBA configured

3.3 Ethernet Port Numbering

In *most cases*, for Windows 2000/2003, the connection named *Local Area Connection* goes to the Ethernet switch in switch bay 1, and the connection named *Local Area Connection 2* goes to the Ethernet switch in switch bay 2. We use the phrase *most cases*, because this is not always the case. For Windows 2000, the order of the Local Area Connection *names* assigned is based on the *order* in which the drivers for each NIC are installed. The drivers necessary to support the NICs on a blade server are not part of a standard Windows 2000/2003 install, and the NICs will be generically listed in Windows 2000 Device Manager as two or more *Ethernet Controllers* until the necessary drivers are loaded. For these NICs to become active, a third-party driver, supplied by IBM, needs to be installed. The *normal* procedure most users follow is to install the drivers on the first *Ethernet Controller* in the list, and then install the drivers on the second *Ethernet Controller* in the list (and so on). The end result of this is the *most-cases* scenario previously mentioned, where the Windows 2000/2003 connection named *Local Area Connection* goes to switch 1 and the one named *Local Area Connection 2* goes to switch 2.

If, however, the drivers are installed on the second *Ethernet Controller* in the list first, and then the first *Ethernet Controller* in the list, the connection names are reversed, and the connection named *Local Area Connection* is now the one going to switch 2, and the connection named *Local Area Connection 2* is now going to switch 1.

In older versions of Linux it was observed that eth0 went to the Ethernet switch in switch bay 2 and eth1 goes to the Ethernet switch in switch bay 1. This was reversed from a *normal* Windows 2000 install, as previously mentioned, and also proves to be the source of some confusion. **However, current testing with Linux (Red Hat Version 3 Update 6) and SLES 9 has shown that eth0 goes to the Ethernet switch in bay 1 and eth1 goes to the switch in bay 2.** On an HS40 blade, 0 & 2 should be connected to top switch module and 1 & 3 connect to the bottom switch module.

The designations can be verified by browsing the MAC address table on the Ethernet switch and match the values to the MAC addresses show in the ifconfig output on the blade. The following example shows the MAC address of Eth 0/1 devices. Log into the top switch module (bay 1) and search for the MAC address on Internal port 13 (Bay 13) corresponding to Eth0 device. It shows the same MAC address of Eth0. For Eth1, do the same by logging into switch module bay2 (bottom) and searching for the MAC address on internal port 13, it also shows the same MAC address for Eth1.

```
[root@hs20-8843 root]# cat /proc/version
Linux version 2.4.21-27.ELsmp (bhcompile@bugs.build.redhat.com) (gcc version 3.2.3
20030502 (Red Hat Linux 3.2.3-47)) #1 SMP Wed Dec 1 21:59:02 EST 2004
[root@hs20-8843 root]# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:11:25:4A:00:4C
          inet addr:9.42.166.82  Bcast:9.42.166.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          ...
```

Browse MAC Address Table – Top Switch

VID	MAC Address	Port	Learned
9	00-11-25-4a-00-4c	Bay 13	Dynamic

```
eth1      Link encap:Ethernet  HWaddr 00:11:25:4A:00:4D
          inet addr:10.0.200.32  Bcast:10.0.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          ...
```

Browse MAC Address Table – Bottom Switch

VID	MAC Address	Port	Learned
100	00-11-25-4a-00-4d	Bay 13	Dynamic

4 Installation Process for Red Hat Linux

4.1 Red Hat Advanced Server 2.1/3.0 Installation

This section is intended to detail the OS installation steps and may reference steps that have been described and completed in previous sections. The steps below assume that the QLogic configuration steps in Section 3.2 have been completed. These steps have been confirmed for Red Hat AS 2.1 Update **4** and Red Hat AS 3.0 Update 3 and above.. Please see Section 4.4 for known issues related to this and other releases of Red Hat.

1. Recommended: From the secondary QLogic Fibre Channel switch module, DISABLE the port of the secondary HBA. (This HBA can be re-enabled after the initial OS load). This process is detailed in 2.4. Note, we have been able to successfully load Red Hat Version 2.1 without this step, but we see "shadow" drives during installation. Please see Section 4.3 for more information relating to dual path installations with Version 3.0.
2. The bootnet.img diskette provided in the Red Hat AS2.1 Update 4 does not contain the network drivers for the network chips on either blades (HS20: tg3; HS40:e1000). A custom bootnet image file is required for using this kickstart process. A single image can be built that works for both HS20 and HS40 blades.
3. One of the differences between installing 2.1 and 3.0 on the blades is that in Version 3.0, the network drivers do not fit on single floppy image with the installation kernel. What we've done to support Kickstart from a floppy is create separate boot diskettes for the HS20 and HS40.
4. Using the Management Module, boot the blade server from the kickstart floppy using the following steps.
   a. First, it is recommended that the media tray is not owned by the target blade. (See Figure 25)

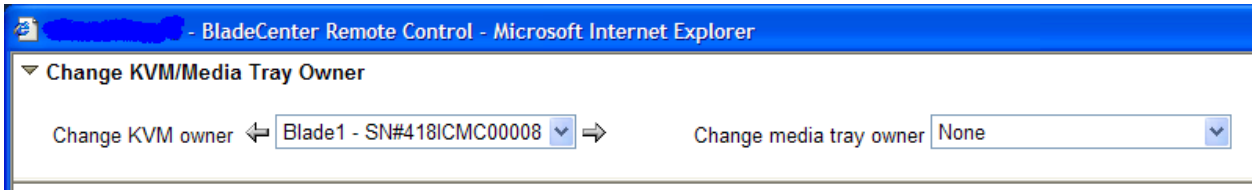

Figure 25 Change Media Tray Owner

   b. Load the kickstart floppy image to the RSA card. This can be done either from a physical floppy disk by selecting "Removable Drive (A:)" or by selecting "Select File…" and specifying a diskette image file. Select one from the left list box, then click on the ">>". (Figure 26) After clicking the ">>", you'll see the dialog shown in Figure 27)

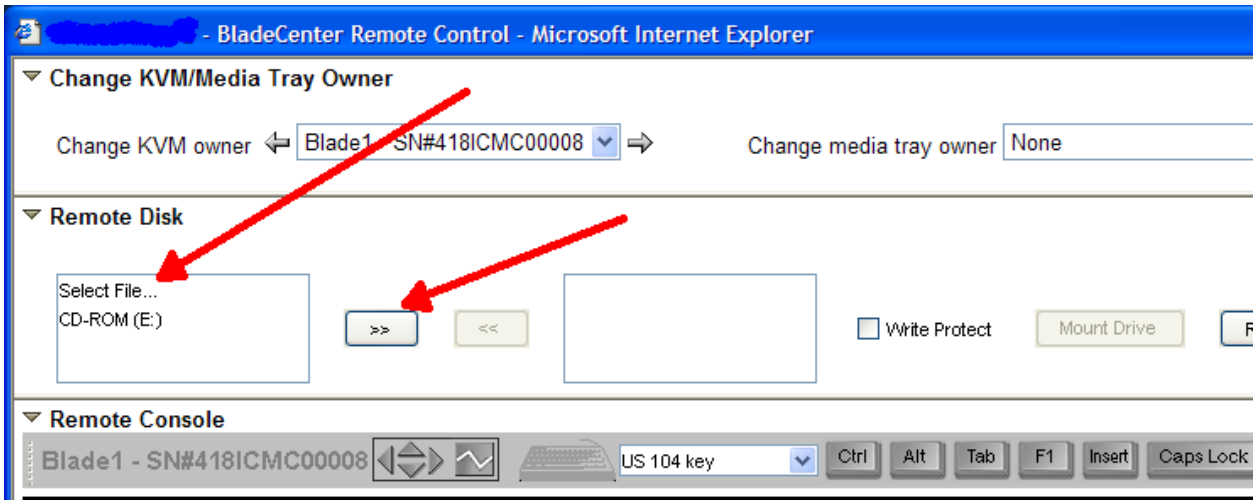

Figure 26 Remote Media via the Remote Console

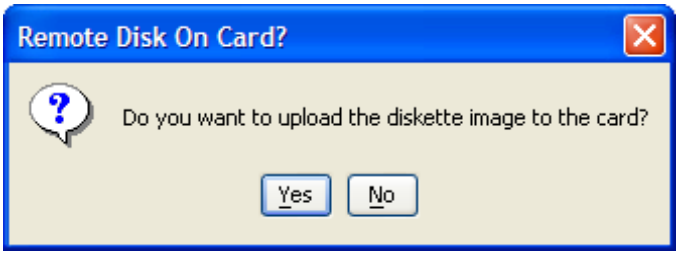

Figure 27 Loading the Remote Diskette

c. Power on or restart the target blade.

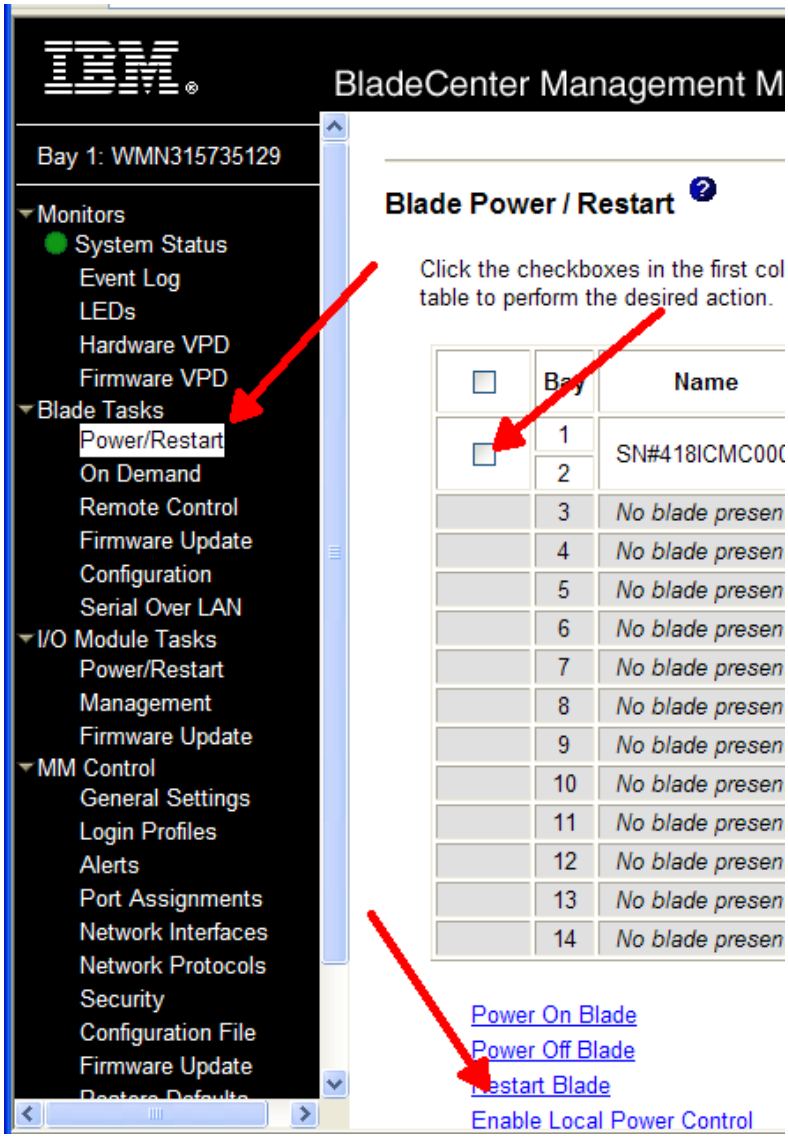

Figure 28 Power On/Restart Blade

5. From this point on, the Red Hat installation proceeds like a standard installation.

Note: you must specify **"linux nofb"**, when installing via BladeCenter Remote Control, otherwise you will see the "No Video Available" screen instead of the Linux installation dialogs. You should also be able to set the correct mode setting with the "vga=0x301" option. See RETAIN Tip 167679 (http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-55660) for more information.

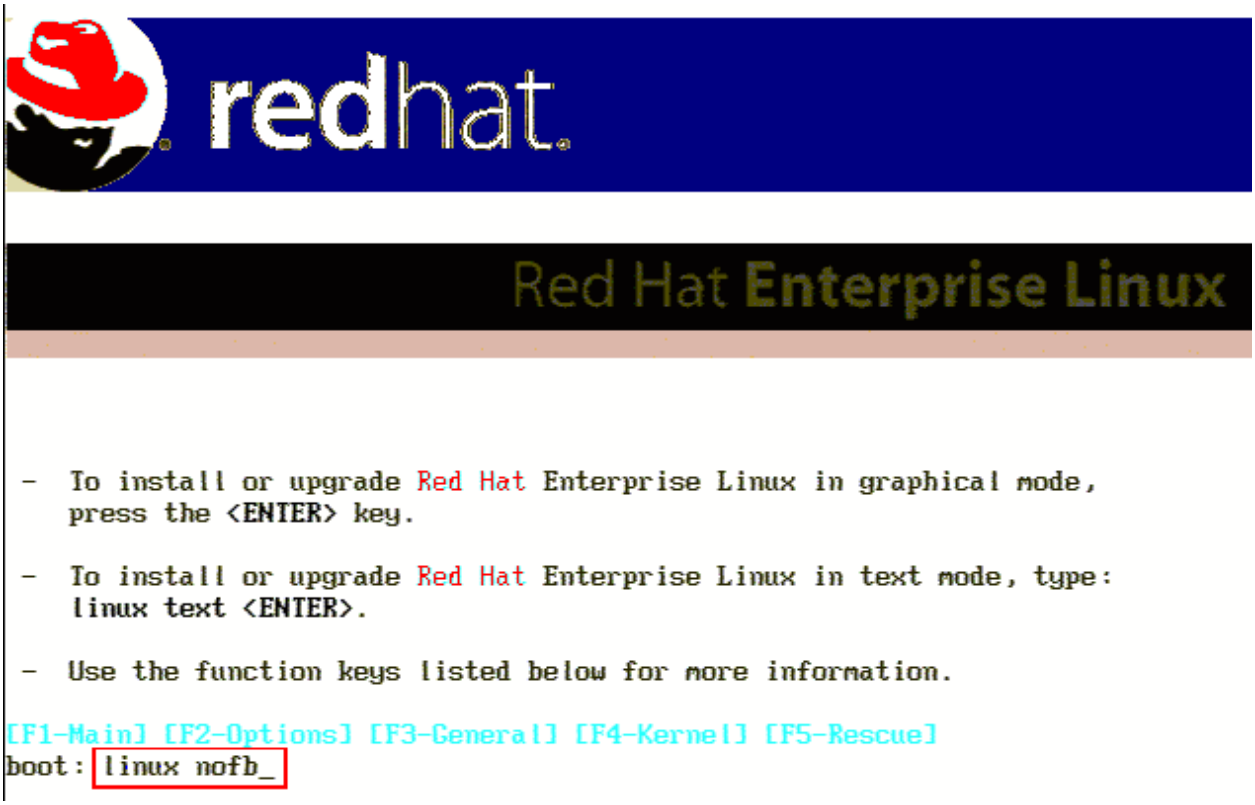

Figure 29 linux nofb

6. **Important:** When the installation is complete, unmount the diskette drive before the reboot.

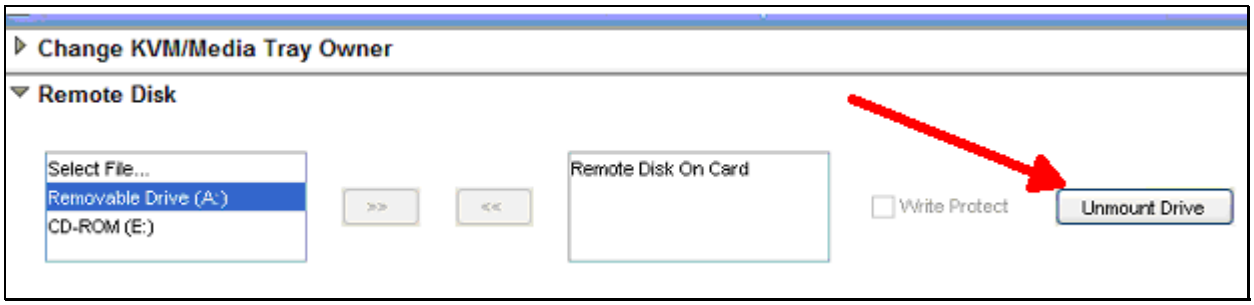

Figure 30 Unmount the Diskette Image

7. When loading Linux on an HS40, after unmounting the diskette, it is possible that EFI may replace the “IBM Remote Disk” in the boot order, and the blade may boot into EFI instead of Linux. If this occurs, follow the steps shown on page 17 to reorder the boot sequence.
8. Depending on your kickstart file configuration, the network interface adapters may not be configured during install. However, you should reference your OS installation guide to configure your respective network interface adapters.
9. Apply the recommended errata and driver updates. If this includes a kernel update, you will need to compile new drivers to comply with the new kernel. (Note: When building the new kernel and driver image please add a new option to the lilo.conf file so that you can revert back to the installed kernel if there were problems with the new driver)
10. Power down the blade server.
11. Once the OS has been installed, you can re-enable the port for the secondary HBA on the Qlogic switch.
12. Zone the McData switch for the WWPN of the second HBA

13. Power on the server and enter the QLogic BIOS by pressing **CTRL+Q**.
14. From the ***Configurable Boot Settings***, set the *BOOT Device* to the alternate controller's World Wide Name (that you gathered) and the *BOOT LUN* to **0**.
15. From the *Basic Settings* panel, verify that the BIOS on this adapter is ***Enabled***.
16. Verify that your storage partitioning and host group configuration looks similar. This configuration should account for both HBA worldwide port names, and the boot partition being configured to LUN **0**.
17. Reboot the server.

4.2 Kickstart Information

The SYSLINUX.CFG should have the following line at the top of the file. Of note: the ks.cfg file is built into initrd.img. "`text`" is specified to avoid problems displaying the graphical installation dialogs via the Management Module that result in blank screens. "`nofb`" is another way to avoid a black screen during install.

```
default ksinstall
      prompt 1
      timeout 600
      …
      label ksinstall
       kernel vmlinuz
       append ks=file://ks.cfg text initrd=initrd.img lang=
                                          devfs=nomount ramdisk_size=7168
```

The kickstart config file will have some blade specific entries:

```
mouse generic3usb --device input/mice
```

Sample partition statements.

```
clearpart –all
part /boot --fstype ext3 --size=128 --ondisk sda --asprimary
part / --fstype ext3 --size=20000 --grow --ondisk sda
part swap --size=2048 --maxsize=2048 --ondisk sda
```

4.3 Multipath Red Hat Installation

While it is possible to install RHAS2.1 with both paths enabled using the standard partitioning statements in the kickstart configuration file, this approach does not work with Version 3.0. The installation will complete and Linux will begin to boot, but you will eventually witness the following error:

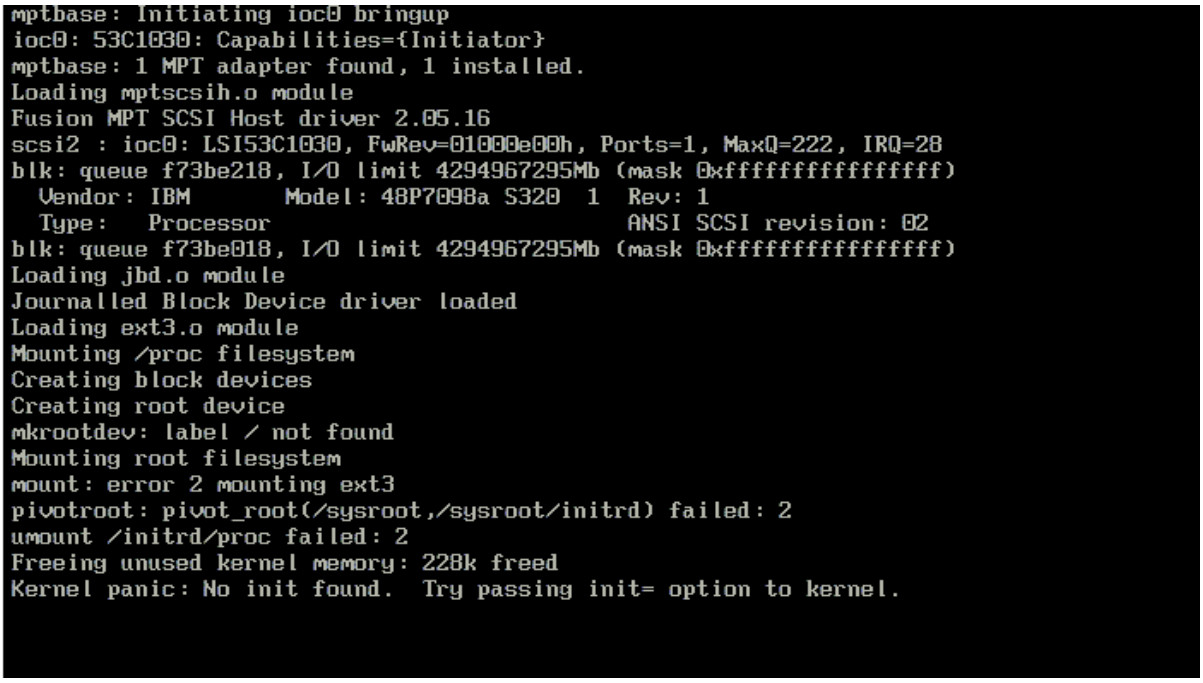

Figure 31 Boot Failure for RH 3.0

Attempting a linux rescue does not appear to work; there are no partitions found on /dev/sda. The problem appears to be with the autopartitioning code and/or the partitioning statements in the ks.cfg. If you are stuck and can't get the second path disabled on the switch, then the following approach *should* work.

The actual partitioning of the drive must be done in the %pre section of the kickstart configuration file, shown in

Figure 32. The partition information is specified using MB values, indicated by the –uM flag. Each partition is defined on a separate "line" with each separated by the \n character.

The actual partitions can be assigned as before, but the sizes are not specified since the partitions have already been created in the pre script.

Do not include the clearpart statement and do not specify zerombr in the kickstart file either. 'zerombr yes' indicates that you would like the Red Hat installation program to initialize all invalid partition tables on disks visible at the time of installation.

```
#zerombr yes
#clearpart --all
part /boot --fstype "ext3" --onpart sda1
part / --fstype "ext3" --onpart sda2
part swap --onpart sda3
%pre
printf '1,100,L,*\n,12000,L\n,2048,S\n,,L\n' | sfdisk /dev/sda -uM
```

Figure 32 Kickstart Modifications

The above example creates the partition information shown in Figure 33.


```
[root@fafrin root]# fdisk -l

Disk /dev/sda: 18.1 GB, 18158714880 bytes
255 heads, 63 sectors/track, 2207 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/sda1   *         1        13    104422   83  Linux
/dev/sda2            14      1543  12289725   83  Linux
/dev/sda3          1544      1805   2104515   82  Linux swap
/dev/sda4          1806      2207   3229065   83  Linux
```

Figure 33 fdisk Output

Note: If you see this message shown in Figure 34, answer no. In this situation, it may be possible to install successfully without the %pre script since it won't attempt to partition the second path.

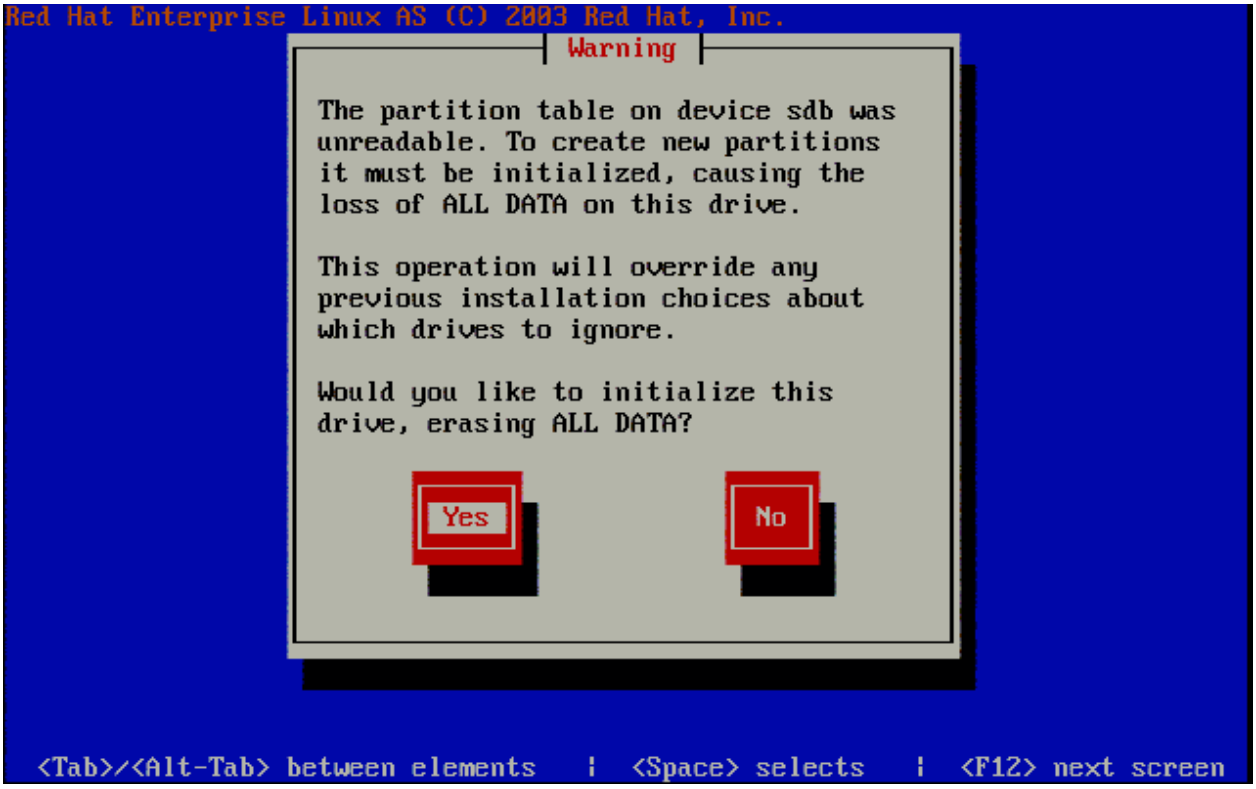

Figure 34 Invalid Partition Table

Aside from the 2 differences noted above, the process detailed in Section 4 is essentially unchanged; just make sure you are using the correct diskette image.

4.4 Red Hat Linux AS Broadcom driver installation

After loading the operating system, the network drivers should work immediately. However, we downloaded the latest Broadcom device drivers for Linux and performed the instructions in Example 6-1 to install them. The latest Broadcom device drivers were obtained from the following URL: http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54186

```
[root@localhost root]# mount /dev/sda /mnt/floppy
[root@localhost root]# ls -al /mnt/floppy
total 285
-rwxr-xr-x 1 root root 279887 Jan 6 20:06 basplnx-6.2.1-1.src.i386.rpm
[root@localhost root]#
[root@localhost root]# cp /mnt/floppy/basplnx-6.2.1-1.src.i386.rpm /tmp
[root@localhost root]#
[root@localhost root]# cd /tmp
[root@localhost tmp]# rpm -ivh basplnx-6.2.1-1.src.i386.rpm
1:basplnx ########################################### [100%]
[root@localhost tmp]#
[root@localhost tmp]# cd /usr/src/redhat/
[root@localhost redhat]#
[root@localhost redhat]# rpmbuild -bb ./SPECS/basplnx.spec
Executing(%prep): /bin/sh -e /var/tmp/rpm-tmp.4397
...
Executing(%build): /bin/sh -e /var/tmp/rpm-tmp.4397
gcc      -DLINUX      -D__KERNEL__      -DMODULE      -I/lib/modules/2.4.9-
e.24smp/build/include -Wall -Wstrict-prototypes -O2 -c blf.c
gcc      -DLINUX      -D__KERNEL__      -DMODULE      -I/lib/modules/2.4.9-
e.24smp/build/include -Wall -Wstrict-prototypes -O2 -c pal.c
ld -r -o basp.o blf.o pal.o blfcore.o
Executing(%install): /bin/sh -e /var/tmp/rpm-tmp.67022
...
Processing files: basplnx-6.2.1-1
Executing(%doc): /bin/sh -e /var/tmp/rpm-tmp.67022
Finding Provides: (using /usr/lib/rpm/find-provides)...
Finding Requires: (using /usr/lib/rpm/find-requires)...
PreReq:  /bin/sh  /bin/sh  /bin/sh  rpmlib(PayloadFilesHavePrefix)  <=  4.0-1
rpmlib(CompressedFileNames) <= 3.0.4-1
Requires(interp): /bin/sh /bin/sh /bin/sh
Requires(rpmlib):       rpmlib(PayloadFilesHavePrefix)       <=       4.0-1
rpmlib(CompressedFileNames) <= 3.0.4-1
Requires(post): /bin/sh
Requires(preun): /bin/sh
Requires(postun): /bin/sh
Requires:     ld-linux.so.2     libc.so.6     /bin/sh     libc.so.6(GLIBC_2.0)
libc.so.6(GLIBC_2.1) libc.so.6(GLIBC_2.1.3)
Wrote: /usr/src/redhat/RPMS/i386/basplnx-6.2.1-1.i386.rpm
Executing(%clean): /bin/sh -e /var/tmp/rpm-tmp.3526

[root@localhost redhat]# rpm -ivh RPMS/i386/basplnx-6.2.1-1.i386.rpm
Preparing... ########################################### [100%]
1:basplnx ########################################### [100%]
[root@localhost redhat]#
```

Figure 35 RPM Install of Broadcom Device Drivers

4.5 Issues related to Red Hat AS 2.1/3.0 and blades

1. The firmware for HS40 blades **MUST** be updated to the minimum levels shown in Table 5 HS40 BIOS and ISMP Information if the blade is going to be running Linux. There is a problem identified in the 1.32 version of the BIOS, related to EFI, which causes processors to be disabled. The only fix for this error is to upgrade the BIOS *and* ISMP firmware.

2. Installing Red Hat by booting from remote CDs is not currently supported on the HS40. Installing the

HS40 from the local media tray will work, but it is recommended that the ***lowres*** option be used for the installation of AS2.1.

3. When booting from the remote floppy, the ***lowres*** (for RHAS2.1) or the ***nofb*** (for RHAS3.0) install mode should function correctly. There have been intermittent instances of no keyboard response.

4. When loading Linux on an HS40, after unmounting the diskette, it is possible that EFI may replace the "IBM Remote Disk" in the boot order, and the blade may boot into EFI instead of Linux. There is also a problem referenced in RETAIN tip H1822293 that results in the HS40 boot order being reordered. (http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-56325) If this occurs, follow the steps shown on page 17 to reorder the boot sequence.

5. When a blade boots and the media tray is assigned to during boot up, the floppy drive in the media tray is assigned as (in this example) /dev/sdc. However, since there is no diskette in the drive, we see the following errors (beginning at READ CAPACITY FAILED):

```
Nov 5 09:44:30 spencerooni kernel: hub.c: new USB device 00:0f.2-1.1, assigned address 3
Nov 5 09:44:30 spencerooni kernel: usb.c: USB device 3 (vend/prod 0x644/0x0) is not claimed by
any active driver.
Nov 5 09:44:30 spencerooni kernel: hub.c: new USB device 00:0f.2-1.3, assigned address 4
Nov 5 09:44:30 spencerooni kernel: usb.c: USB device 4 (vend/prod 0x5ab/0x30) is not claimed
by any active driver.
Nov 5 09:44:30 spencerooni kernel: Initializing USB Mass Storage driver...
Nov 5 09:44:30 spencerooni kernel: usb.c: registered new driver usb-storage
Nov 5 09:44:30 spencerooni kernel: scsi1 : SCSI emulation for USB Mass Storage devices
Nov 5 09:44:30 spencerooni kernel:   Vendor: TEAC      Model: FD-05PUB          Rev: 2000
Nov 5 09:44:30 spencerooni kernel:   Type: Direct-Access   ANSI SCSI revision: 02
Nov 5 09:44:30 spencerooni kernel: Attached scsi removable disk sdc at scsi1, channel 0, id 0,
lun 0
Nov 5 09:44:30 spencerooni kernel: sdc : READ CAPACITY failed.
Nov 5 09:44:30 spencerooni kernel: sdc : status = 1, message = 00, host = 0, driver = 08
Nov 5 09:44:30 spencerooni kernel: Current sd00:00: sense key Not Ready
Nov 5 09:44:30 spencerooni kernel: Additional sense indicates Medium not present
Nov 5 09:44:30 spencerooni kernel: sdc : block size assumed to be 512 bytes, disk size 1GB.
Nov 5 09:44:30 spencerooni kernel:  sdc: I/O error: dev 08:20, sector 0
Nov 5 09:44:30 spencerooni kernel:  I/O error: dev 08:20, sector 0
Nov 5 09:44:30 spencerooni kernel:  unable to read partition table
```

Now, that wouldn't be the end of the world, necessarily, However, ANYTIME the KVM is switched on the Bladecenter to/from that blade, the USB "stuff" (KVM and media tray) get re-hotplugged, for lack of a better term, and the sdc errors are generated. They occur even after switch the media tray owner to a different blade. The result is too many spurious errors which is making the internal support team cranky.

Modification 1: Do not boot blade with media tray assigned in order to avoid errors when someone takes the media tray away AND when the floppy drive is empty. NOTE: "alias floppy off" should be added to modules.conf to avoid floppy drive errors when media tray not assigned on boot.

One final note on this item: currently the media tray is required for installing Linux on the blades, but should only be needed in rare circumstances after that. Additionally, since there appear to be several combinations of events that will cause the EFI shell boot option to be restored in the boot sequence on a blade (see item 4), it is recommended that boot sequence be confirmed when an HS40 is rebooted until the firmware fix is released. This can be done by simply confirming that the HS40 boots from the HBA during a restart. So these items should be included as required steps for building a Linux blade.

Modification 2: Leave read-only floppy diskette in the drive. This prevents the "Medium not present" errors from occurring. This is not a complete solution but is a good start. See #3 for more.

Modification 3: When media tray is assigned to a blade, it will automatically hotplug itself and the cdrom and floppy drive will be available for use. When user is done with the cdrom and/or floppy, media tray should be unassigned from the blade and/or the usb-storage module should be removed. The command for this is "rmmod usb-storage".

This step should also be performed when the IBM Remote Disk is used. The IBM Remote Disk is the feature where an image, or even the local floppy or CDROM drive can be mounted on the blade from the local machine.

I have not seen anything to indicate that a Red Hat kernel fix can be completed to prevent these error messages from being generated.

6. floppy.o error messages. Adding 'alias block-major-2 off' to /etc/modules.conf will get rid of that noise.
7. There is a know bug in anaconda which causes the installation of Red Hat 2.1 Update 3 to fail in most cases on an HS40. (http://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=113346) This is the charming "*You're running me on a live system! That's incredibly stupid.*" message and has been addressed in errata (rhn.redhat.com/errata/RHEA-2004-148.html) which is included in RHAS 2.1 Update 4.
8. An error message of "pc_key_b: controller malfunction (0xA7)" is seen many times when RHEL 3 is booting during a five-second interval. This problem exists because the IBM eServer BladeCenter HS20, type 8843, is the first legacy free x86 system, and therefore, does not have a standard keyboard controller. This enhancement allows the operating system to use of the keyboard controller IRQ for other devices. This error message has no negative impact to the system and can be ignored. (http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-57605)
9. Linux experiences kernel panics or hangs during boot of 8843 blades. Hang typically happens when reaching net4 frame diverter after 'real time clock driver.'

   Fix is to add the **usb-handoff** kernel directive within the bootloader configuration file. This is intended to resolve a problem with, rather than eliminate, usb. Here is a summary of the usb issue:

   When the BIOS is configured with legacy USB support enabled it can generate a system management interrupt (SMI) during the boot sequence. The BIOS routine which handles that action clobbers a CPU register. That, understandably, is a bad thing as the Linux kernel does not expect its data to be randomly changed by asynchronous events like that SMI. Whether or not the corruption is catastrophic depends on the exact timing of events (down to the microsecond). That's why sometimes the bootup is successful.
10. If you suffer major path problems (LIPs) or controller panics, it can hang the server indefinitely as the system attempts to find a stable path.
11. If there is a path failure and the host is generating I/O, the boot Drive will move to the other path. However, while this transition is occurring, the system will appear to freeze for up to 30 seconds.
12. Mouse support for the BladeCenter unit

    The ability of the remote console applet to accurately track the location of the mouse cursor in a Linux X Window System session depends on the configuration of the X Window System. Complete the

procedures in the following sections to configure Red Hat Linux and the X Window System for accurate mouse tracking. Type the commands through the remote console or at the keyboard that is attached to the BladeCenter unit.

1. Providing resolution information to the mouse handler

   The mouse handler (mousedev) must have the correct resolution information to track the remote mouse cursor correctly. If the host (blade server) resolution is not 1024 x 768, complete the following steps to provide the correct resolution data to the mouse handler:

   a. If you need to switch to text mode, type: `init 3`
   b. To unload the mouse device driver module, type: `rmmod mousedev`
   c. To notify the mouse device driver of the video resolution, add the following line to the /etc/modules.conf file:
      `options mousedev xres=x yres=y`
      (where x and y are the video resolution)
   d. To reload the mouse device driver module, type: `insmod mousedev`
   e. To return to GUI mode, if necessary, type: `init 5`

2. Changing the pointer speed settings

   To enable the remote mouse to track properly in Red Hat Linux, set the mouse acceleration and threshold values to 1. To change mouse acceleration and threshold values for the current session, open a terminal window and type: `xset m 1 1`

   Complete the following steps to preserve these changes between sessions. NOTE: A remote mouse is not supported in the text console.

   * If you are using GNOME:
     1. Press Ctrl+Esc.
     2. From the menu, click Programs->Settings->Session->Session properties->Startup programs.
     3. Click the Startup programs tab.
     4. Click Add.
     5. In the command line, type: `xset m 1 1` and click OK.
     6. Click Apply and OK to close the window.
     7. Log out from the session and select “Save current setup” in the Log Out window.

   The next time you log in, the remote and local mouse are synchronized.

   * If you are using KDE:
     1. Using the keyboard, press Alt+F1.
     2. From the menu, click Preferences --> Peripherals --> Mouse.
     3. Click the Advanced tab and change the Pointer Acceleration and Threshold values to 1.
     4. Click Apply and OK to close the window.
     5. Log out from the session and select “Save current setup” in the Log Out window.

   The next time you log in, the remote and local mouse are synchronized.

3. Synchronizing the remote and local mouse for remote control sessions

   If the remote and local mouse are not synchronized correctly, that is, if there is a constant offset between them on the remote client window, move the mouse toward the bottom of the monitor screen in the remote client window. Then, move the local mouse back in the window. This will correct the alignment between the mice.

4. Correcting mouse cursor click related issues

   NOTE: The blade server might not start in GUI mode if the mouse is not detected when the blade server is turned on. This might happen if the blade server is not associated with the keyboard, monitor, and mouse. If there is no mouse cursor or if single clicks from the remote mouse do not work, complete the following steps to correct the problem:

   1. Open a terminal window.
   2. Type setup.
   3. Click Mouse Configuration.
   4. Scroll down and select the generic or two-button USB mouse. Click OK.
   5. Click Yes to update X Configuration.

6. Exit the utility.
7. Press Ctrl+Alt+Backspace.

The next time you log in to the X Window System, the remote mouse will work correctly.

13. Boot-from-SAN fails with Gatekeeper LUN on EMC backed storage – When installing in a boot-from-SAN configuration on an EMC Symmetrix storage array, the following message appears: "The `partition table on device sda was unreachable. To create new partitions it must be initialized.`" Or after rebooting, you simply see a blinking cursor.

This happens because the Symmetrix uses LUN 0 as a gatekeeper LUN, which is marked as a pseudo-device. The Linux installer maps LUN 0 to `/dev/sda`. When the installer tries to configure the boot LUN, it is unable to use `/dev/sda`, so it uses `/dev/sdb` instead. This leads to a problem with identifying the boot loader.

There are two possible solutions to this problem. One solution is to mask or renumber the LUN of the gatekeeper; the other is to change the boot loader configuration.

1. Masking or Renumbering the Gatekeeper (VCM) LUN

To boot from SAN, ESX Server requires that the boot LUN be the lowest-numbered LUN in the storage array. To accomplish this, you may follow one of two approaches (in order of preference):

- To prevent the installer from attempting to write to this device, EMC recommends masking the LUN.
    1. Initialize the Volume Logix database using the administrative host.
    2. Modify the active configuration file to enable the `fba flag2`. This restricts access to the Volume Logix database.
- The gatekeeper (VCM) LUN can be assigned a number higher than the boot LUN number. Current Fibre Channel adapters can boot only from LUN numbers up to 15, so your boot LUN number must be in the range 0 to 15. Renumber the gatekeeper (VCM) LUN to 16 or higher, and it won't interfere with the boot LUN.

To renumber the LUN, either contact the vendor for assistance or consult your storage array management documentation.

2. **Changing the Boot Loader Configuration**

Another way to deal with the gatekeeper LUN is to specify in your boot configuration where to find the boot LUN. This requires editing the `/etc/lilo.conf` file.

Note: If you have already completed the installation and you are unable to boot the server, you should boot from a Linux rescue CD and mount the boot LUN so you can make the following changes.

1. Back up the file `/etc/lilo.conf`.
2. Edit the file `/etc/lilo.conf`.
3. Find the line that says `default=esx`.
4. After the `default` line, insert these two lines:
   disk=/dev/sdb
   bios=0x80

   Use a tab to indent the second line.

5. Exit the editor, saving your changes.
6. Run `lilo`.
7. Reboot the server. This time, it should boot successfully from `/dev/sdb`.

4.6 Other Linux Distributions and Blades

We've just begun our own testing with SLES 9 and Red Hat 4 on blades in a boot to SAN environment and we hope to add more material here in the future. At this point, the only new problem encountered was that our 8832 blade would not boot from the .iso image of SLES9 mounted via the Management Module. It would boot from the actual CDROM mounted remotely, and the install.

If you are installing SLES9 to a local drive on an HS20, watch out for this RETAIN tip: http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-57841.

4.7 RDAC Installation and Configuration

Notes on Linux

There are two versions of the Linux RDAC. The version (09.00.A5.09)for Linux 2.4 kernels only like Redhat EL 3 and SuSe SLES 8 and RDAC package version 09.01.B5.XX for Linux 2.6 kernel environment.

- Make sure you read the README.TXT files for V9.15 Linux RDAC, HBA and Storage Manager for Linux.
- When using the Linux RDAC as the failover/failback driver, the host type should be set to LNXCL instead of Linux. If Linux RDAC is not used, the host type of Linux must be used instead.
- There is not a requirement that the UTM (Access LUN) must be removed from the LNXCL Storage Partitioning partition.
- The Linux RDAC driver cannot co-exist with a HBA-level multi-path failover/failback driver such as the 6.06.63-fo driver. You might have to modify the driver make file for it to be compiled in the non-failover mode.
- Auto Logical Drive Transfer (ADT/AVT) mode is not supported when using Linux in a cluster, or when using RDAC. Since ADT(AVT) is automatically enabled in the Linux storage partitioning host type. It has to be disabled by selecting 'host type' of LNXCL.
- AVT is required if you are using the Qlogic Failover drivers.
- The Linux SCSI layer does not support skipped (sparse) LUNs. If the mapped LUNs are not contiguous, the Linux kernel will not scan the rest of the LUNs. Therefore, the LUNs after the skipped LUN will not be available to the host server. The LUNs should always be mapped using consecutive LUN numbers
- Although the host server can have different FC HBAs from multiple vendors or different FC HBA models from the same vendors, only one model of FC HBAs can be connected to the IBM DS4000 Storage Servers.
- If a host server has multiple HBA ports and each HBA port sees both controllers (via an un-zoned switch), the Linux RDAC driver may return I/O errors during controller failover.
- Linux SCSI device names have the possibility of changing when the host system reboots. We recommend using a utility such as **devlabel** to create user-defined device names that will map devices based on a unique identifier, called a UUID. The **devlabel** utility is available as part of the Red Hat Enterprise Linux 3 distribution, or online at: http://www.lerhaupt.com/devlabel/devlabel.html
- Linux RDAC supports re-scanning to recognize a newly mapped LUN without rebooting the server. The utility program is packed with the Linux RDAC driver. It can be invoked by using either **hot_add** or **mppBusRescan** command (note that **hot_add** is a symbolic link to **mppBusRescan**). There are man pages for both commands. However, the Linux RDAC driver doesn't support LUN deletion. One has to reboot the server after deleting the mapped logical drives

4.8 Serial Over LAN (SOL) for Blades

4.8.1 Overview

The Management Module command-line interface provides access to the text-console command prompt on each blade server through a Serial over LAN (SOL) connection, enabling the blade servers to be managed from a remote location.

To start an SOL connection with a blade server, you must first start a telnet or ssh session with the management module. When this session is running, you can start a remote-console SOL session with any blade server in the BladeCenter unit that is set up and enabled for SOL operation. You can establish up to 20 separate Telnet sessions with a BladeCenter management module. For a BladeCenter unit, this enables you to have 14 simultaneous SOL sessions active (one for each of up to 14 blade servers) with 6 additional command-line interface sessions available for BladeCenter unit management.

There are several well written white papers available that describe the steps required to configure Serial Over LAN for the BladeCenter. The firmware sections of these documents can be skipped as long as you have followed the firmware recommendations in this document. For Linux, you must make the changes to both /etc/inittab and your boot loader configuration file (either lilo.conf or grub.conf) in order to get the full functionality under Linux.

4.8.2 Summary of Linux changes for SOL

1. Add the following line to the end of the # Run gettys in standard runlevels section of the /etc/inittab file. This enables hardware flow control and enables users to log in through the SOL console.
   `7:2345:respawn:/sbin/agetty -h ttyS1 19200 vt102`
2. Add the following line at the bottom of the /etc/securetty file to enable a user to log in as the root user through the SOL console:
   `ttyS1`
3. LILO configuration
   In order to enable SOL in Linux, we need to add 2 parms to the kernel line in the bootloader configuration file:
   `console=ttyS1,19200n8`
   `console=tty1`
   The order of these two parms affects where a portion of the Linux boot (and shutdown messages) are displayed, and consequently how the Linux boot can be managed. Basically, your two options are having the messages go to the main console (the one you see via Remote Control) or the serial console (SOL). The existing documentation refers to these options as being able to monitor the boot via SOL, or being able to interact with the boot via SOL.

   If you are using LILO, complete the following steps:
   1. Complete the following steps to modify the /etc/lilo.conf file:
      a. Add the following text to the end of the first `default=linux` line
         `-Monitor`
      b. Comment out the `message=/boot/message` line by adding a # at the beginning of this line. This disables the graphical LILO boot screen since it will not display correctly in the SOL session. The next time the blade is booted, the text LILO: prompt appears in place of the graphical boot screen.
      c. Increase the timeout value to give the operator a better chance to choose which way the blade boots.
      d. Add the following line before the first image=... line:
         `# This will allow you to only Monitor the OS boot via SOL`
      e. Add the following text to the end of the first label=linux line
         `-Monitor`
      f. Add the following line to the first image=... section. This enables SOL.
         `append="console=ttyS1,19200n8 console=tty1"`
      g. Add the following lines between the two image=... sections:
         ```
         # This will allow you to Interact with the OS boot via SOL
         image=/boot/vmlinuz-2.4.21-37.ELsmp
         label=linux-Interact
         initrd=/boot/initrd-2.4.21-37.ELsmp.img read-only
         root=/dev/hda6 append="console=tty1 console=ttyS1,19200n8"
         ```
   2. Run the lilo command to store and activate the LILO configuration: `/sbin/lilo -v`.

   When the Linux operating system starts, a LILO boot: prompt is displayed instead of the graphical user interface. Pressing Tab while at this prompt will install all of the boot options that are listed. To load the operating system in interactive mode, type linux-Interact and then press Enter.

To see the choices of kernels for boot from, simply hit the tab key.To re-enable the graphical boot screen, replace or uncomment the above line in the `lilo.conf` file and rerunning LILO.

4.8.3 alt-sysrq

One of the reasons you may be configuring SOL is for utilizing the alt-sysrq “magic key sequence”.

Enter the command:

```
> sysctl -w kernel.sysrq = 1
```

This will update that value in `/etc/sysctl.conf` and also set the value in `/proc/sys/kernel/sysrq` to 1.

4.8.3.1 SOL using a Linux client

1. First, telnet to the MM and login. Then, enter the command

   `console -T system:blade[x]`

   where x is the target blade number.
2. (here is the trick) Once you have an SOL connection, hit the escape sequence which is **`ctrl-]`.** This should bump you to a telnet prompt.
3. Next, type **`send brk`** and hit enter. You'll get a new line, but no output.
4. Then hit the sysrq command key you want such as h or m or p or w.

It should look like this:

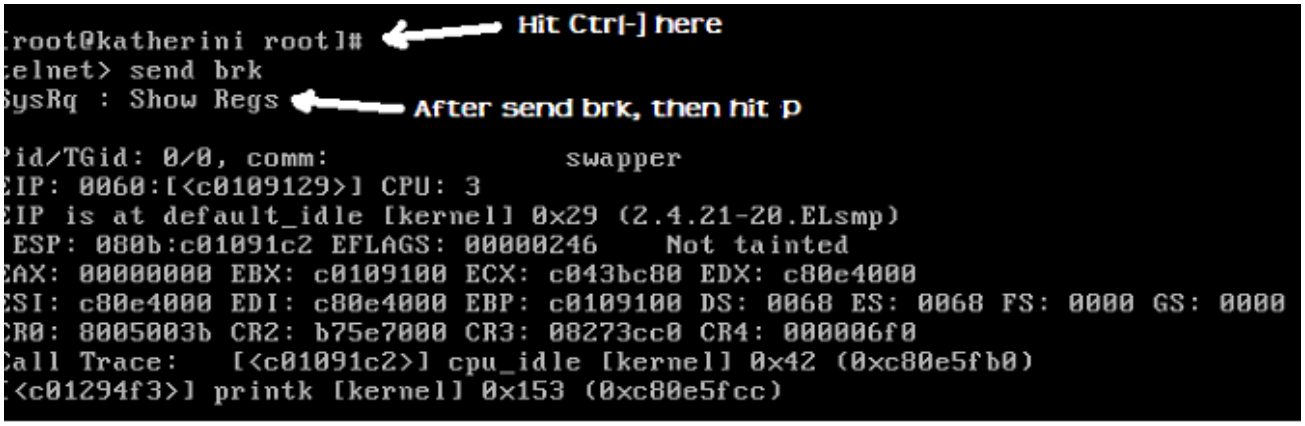

Figure 36 alt-sysrq from Linux client

4.8.3.2 alt-sysrq from a Windows Client

For a Windows client, you need to use PuTTY, or alternately, the telnet client from the cygwin package. I was not able to get this working using regular telnet application in Windows. Hyperterminal didn't work either.

1. First, telnet to the MM and login. Then, enter the command

   console -T system:blade[x]

   where x is the target blade number.
2. (here is the trick) From the PuTTY pulldown menu, click on Special Command -> Break.
3. Then hit the sysrq command key you want. Start with `sysrq h.` Use `sysrq 8` for maximum output level.

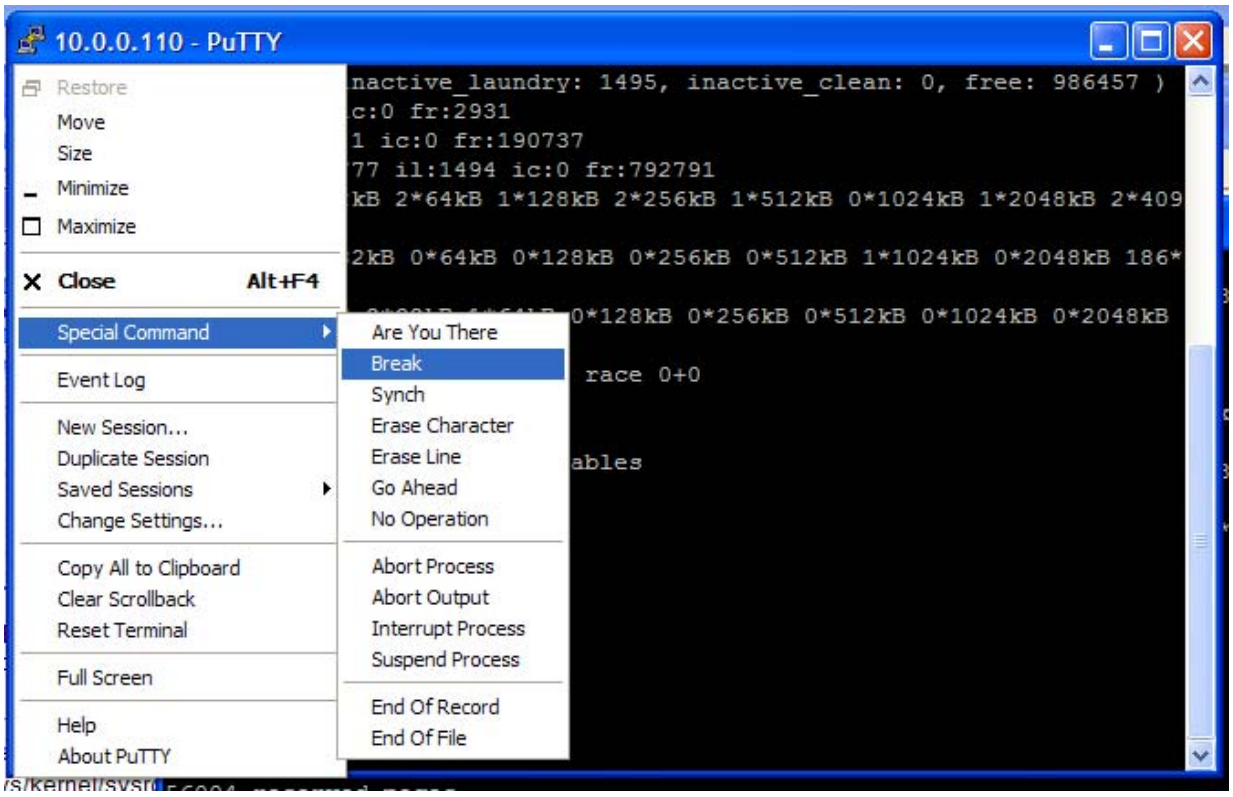

Figure 37 alt-sysrq from PuTTY

Note, PuTTY has a nice feature that will log all of your session output to a file. Specify that before opening the telnet session.

In certain rare circumstances, the serial data of the 8843 Blade Server may stop flowing. This results in a blank SOL session with only a blinking cursor. When this occurs, simply press the control and q keys (Ctrl+q) at the same time. This key sequence will allow the flow of serial data to resume.

Notes: Telnet sessions with the BladeCenter management module have a default timeout value of 120 seconds. This means the telnet session will remain active indefinitely. This value may be customized to a selectable value within the management module. I would recommend changing the telnet/SOL timeout to "no timeout". Default is 120 seconds. This can be changed via the web interface, under MM Control->Network Protocols->Telnet Protocol or via the CLI:

```
system> telnetcfg ?
telnetcfg {-t {<secs>}}
Displays and configures the telnet settings for the MM
 -t: inactivity timeout in secs (valid value from 0 to 4,294,967,295 seconds)
```

This command sets it to no timeout: (t=0)

```
system> telnetcfg -t 0 -T system:mm[1]
```

Although this says telnet config, it works for both telnet and SOL.

4.9 Netdump

Netdump is a Red Hat specific crash dump facility. Unlike traditional crash dump facilities, this facility dumps memory images to a centralized server via the network. More information on netdump can be found here:

http://www.redhat.com/support/wpapers/redhat/netdump/

Here is the link to the instructions from Red Hat that can be used to set up netdump:
http://www.redhat.com/support/wpapers/redhat/netdump/setup.html

The netdump server from the RHEL3U3 distribution is version 0.6.11-3.

Basically, here is what needs to be done at the server:

1) rpm -i netdump-server-0.6.11-3.i386.rpm
2) passwd netdump
3) chkconfig netdump-server on
4) service netdump-server start

And at the client

1) netdump client package probably already installed.
2) service netdump propagate
3) modify /etc/sysconfig/netdump. Add line NETDUMPADDR=10.0.0.26
4) chkconfig netdump on
5) service netdump start

Note, there have been instances reported where there have been dependency issues between netdump and openssh. Here are the openssh packages that are installed on our test blade:

```
[root@katherini tmp]# rpm -qa | grep ssh
openssh-clients-3.6.1p2-33.30.1
openssh-3.6.1p2-33.30.1
openssh-server-3.6.1p2-33.30.1
```

5 Install Process for Windows 2000 and 2003 Remote Boot

5.1 Operating System Installation

1. Obtain the Fibre Daughter card device driver specified by your storage vendor to support the use of the Qlogic HBA with their storage. Take the installation package and extract files and place the directory structure onto a driver diskette.
2. Boot the blade and access the Windows installation media.
3. As the Windows installation media is accessed, a message will appear prompting the user to press the “F6” function key to load a device driver.
4. Depress the key and when prompted, load the HBA device driver created in step 1
5. **If Windows prompts you to use the driver included with Windows or the driver that you have specified, be certain to use the driver from the diskette, not the driver included with Windows.**
6. Continue with the Windows OS load as normal.
7. Install the recommended service packs and reboot.
8. For path failover, a multi-path device driver must be installed as specified by the storage vendor.
9. Power down the HS20 blade server.
10. Once the OS has been installed, you can re-enable the port for the secondary HBA on the Qlogic switch.
11. If required, zone the external switch for the WWPN of the second HBA
12. Power on the server and enter the QLogic BIOS by pressing **CTRL+Q**.
13. (For port B – 2600) From the ***Configurable Boot Settings***, set the *BOOT Device* to the alternate controller's World Wide Name (that you gathered) and the *BOOT LUN* to **0**.
14. From the *Basic Settings* panel, verify that the BIOS on this adapter is ***Enabled***.
15. Verify that your storage partitioning and host group configuration looks similar. This configuration should account for both HBA WWPNs, and the boot partition being configured to LUN **0**.
16. Reboot the server, setup is complete.

5.2 Installing Broadcom NetXtreme Gigabit Ethernet device drivers

Complete the following steps to install Broadcom NetXtreme Gigabit Ethernet device drivers:

1. Insert the CD or diskette that contains the Broadcom NetXtreme Gigabit Ethernet device drivers into the appropriate drive. The CD or diskette contains an InstallShield wizard that installs the device drivers for the device. For information about downloading the Broadcom NetXtreme Gigabit Ethernet application, see 2.0 Where to download device drivers and files.

NOTE: Steps 2 through 5 can be skipped in this section if there will not be a need to extract files. However, if the driver was downloaded from the IBM website, the downloaded file must be extracted.

2. Select the application. In the InstallShield wizard window, select Next.
3. Enter the location of the folder where the files are to be saved. Make note of this location for future reference. Later, you will need to point to this location when installing your device drivers. Select Next.
4. The InstallShield wizard extracts the files needed to install all Broadcom Gigabit Ethernet device drivers on the blade server. When InstallShield is finished, select Next.
5. On the desktop, right-click My Computer and select Manage. Select Device Manager. The Broadcom NetXtreme Gigabit Ethernet device is detected automatically as an unknown device under Other Devices. Right-click one of the unknown Ethernet devices. Select Properties -> Reinstall Driver. In the Welcome to Upgrade Device Driver Wizard window, select Next.
6. Select Search for a suitable driver for my device and select Next.
7. From the Optional Search Location menu, select Specify a location, and select Next.

8. Type the location of the Ethernet device drivers that you extracted in step 6. Select OK.
9. When the message "System found a driver for device" is displayed, select Next to install the device driver.
10. When you are prompted to restart the blade server, select Yes if you do not have more device drivers to install, or select No if you have another device driver to install.

5.3 Broadcom Advanced Control Suite installation

Network interface card (NIC) teaming is one method for providing high availability and fault tolerance in IBM Eserver servers. In this example, we use Broadcom Advanced Server Program (BASP) to implement teaming functionality along with load balancing, fault tolerance, and VLAN tagging.

To enable NIC teaming, the Broadcom Advanced Control Suite application must be used on the HS20s. The program is included with the drivers, which you can download at: http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-43815.

To install the suite, perform the following steps:

1. Navigate to the location where the Broadcom Advanced Control Suite application files were extracted (default C:\Drivers\BcomXXX, where XXX is the code level). Execute Launch.exe.
2. Click MANAGEMENT PROGRAMS, and a window similar to the one in Figure 45 opens.

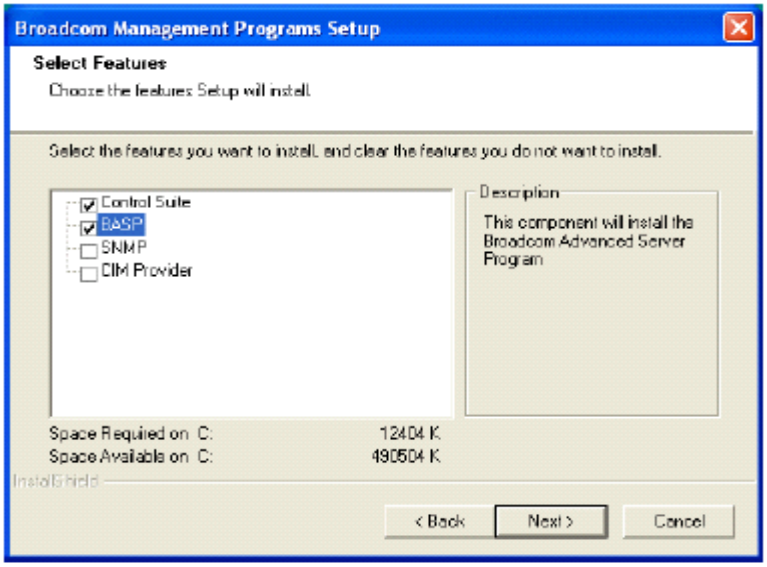

Figure 38 Broadcom Setup

3. Select Control Suite and BASP. Click Next to continue and click then **Finish.**

5.4 List of limitations

1. If there is a path failure and the host is generating I/O, the boot drive will move to the other path. However, while this transition is occurring, the system will appear to freeze for up to 30 seconds.
2. You cannot enable the BIOS for both adapters at the same time. If you do and there is a path failure on the primary adapter (and the adapter is still active), the system will trap with an INACCESSIBLE_BOOT_DEVICE error on reboot.
3. If the boot device, (LUN 0), is not on the same path as the bootable HBA port, you will receive an INACCESSIBLE_BOOT_DEVICE error message.
4. If you suffer major path problems (LIPs) or controller panics, it can hang the server indefinitely as RDAC tries to find a stable path.
5. The IDE disk devices should not be re-enabled.

6 Remote Update of the Firmware on an HS20 or HS40

The steps below detail how to update the blade firmware remotely. This process is much quicker than updating the firmware locally using the media tray since the slow diskette reads are replaced by much quicker "virtual floppy" reads from the Management Module. The downside is we are exposed to some of the network related management module quirks we have experienced.

6.1 BIOS and ISMP

1. Go here to download the most recent versions of the BIOS and ISMP updates for the HS40. Currently there are three files that you'll need to download: 2 for the BIOS update and 1 for the ISMP update. (www-306.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54597)
2. The files are downloaded as "self-extracting" disk images. Perform the following steps for each: Extract the diskette image to a real floppy disk, then create a disk image file of the physical floppy disk using a program such as DSKIMAGE. Save the images with easily identifiable names.
3. Next, upload the first BIOS diskette image to the Management Module. See Figure 26 Remote Media via the Remote Console for an example of how to perform this.
4. Once the diskette image has been loaded, restart the target blade.
5. The blade will boot to DOS. You'll be asked several questions; the default answers are acceptable.
6. After several minutes of reading, erasing, writing, and verifying, the update program will ask for the second BIOS diskette. Unmount the media (Figure 30) and then upload the second BIOS diskette to the management module. After the upload completes, return focus to the remote console window and hit enter.
7. When the BIOS update completes, it will ask you if you'd like to reboot. Before answering yes, unmount the second BIOS diskette, and upload the ISMP diskette. Then ok the reboot.
8. The ISMP diskette will boot to a DOS prompt. At the DOS prompt, enter **LCREFLSH.**
9. When the ISMP update is finished, you will be returned to a DOS prompt. Unmount the remote disk, and reboot.

6.2 BMC Update for the HS20 (8843) and LS20

BMC firmware update for the 8843 blade can be done either by using a bootable diskette image available from ibm.com, or remotely by using PKT files and applying them via the Management Module firmware update interface. While this flashing via the Management Module apparently can be done while the blade is running, the final apply of the firmware WILL cause the blade to power off. So the target blade should be power off AND all blades must be discovered before flashing remotely.

To install the BladeCenter BMC firmware from the BladeCenter Management Module web interface, perform the following steps:

1. In the navigation frame, click Firmware Update under Blade Tasks, select the Target blade slot number, then click Browse.
2. Switch to the drive and directory containing the unpackaged files. The package files are found on Diskette 2. Navigate to the PKT file you want to update, named as:
   BWBTxxy.PKT (Where xx = build ID, and y = revision.)
3. Click Open. The file (including the full path) appears in the box beside Browse.
4. To begin the update process, click Update. A progress indicator opens as the file is transferred to temporary storage. A confirmation page opens when the file transfer is completed and the BladeCenter Integrated System Management Processor has been restarted.

In case of a failed RS485(remote) flash

1. Only one blade can be in kernel mode in a chassis at one time.
2. The Management Module has known issues if a flash fails. The easiest way to recover is to
   a. Pull the failed blade from the chassis
   b. Reset the MM
   c. Reinsert the failed blade
   d. Retry the flash on the failed blade after the blades have all been discovered.

Notes:

1. It is necessary to upgrade the BIOS on the HS40 to the 1.44 level and the ISMP to 1.24 if you are installing Linux on the blade. (See Table 5) There is a problem identified in the 1.32 version of the BIOS, related to EFI, which causes processors to be disabled. The only fix for this error is to upgrade the BIOS *and* ISMP firmware.
2. Occasionally the boot of the first BIOS diskette will fail with old time DOS messages such as “cannot find country.sys.” If you end up at a DOS prompt when booting the BIOS diskette, simply run AUTOEXEC.BAT to kick off the update process. If you don’t get a DOS prompt, restart the blade.
3. Take care when mounting the second BIOS diskette and try to avoid having to restart the Management Module once the update process has begun. If you do encounter and unrecoverable error at this point, you should be able to restart the BIOS process from the beginning. The blade will most likely boot off of the backup BIOS. In this case, answer ‘no’ when asked if you’d like to backup the primary BIOS.
4. If these steps are performed after Linux has been installed, it is advisable to confirm that the blade reboots back to Linux and not into EFI.

7 Flash the QLogic HBA BIOS

Perform the following steps to flash the QLogic HBA BIOS:

1. Create a blank, bootable disk formatted with MS-DOS, PC-DOS, or DR-DOS.
2. Using a text editor, verify that the config.sys and autoexec.bat files are blank and contain no entries.
3. Insert the disk created in step 1 in an appropriate drive. Double-click the archive icon, then type the following in the **Unzip to folder** field to extract the contents of the BIOS files archived to the disk: `a:\`.
4. Insert the disk from step 3 into the drive mounted on the target blade and boot from the disk.
5. At the a:\ prompt, type `ibmutil /i /f`, then press ENTER. This will update the BIOS for both channels of the expansion card in the system.
6. If the QLogic HBA is detected, the following lines display:

```
Current version of the QLA2312 flash BIOS at I/O Address 2400 is 1.38
New version of the QLA2312 flash BIOS at I/O Address 2400 is 1.43
Programming QLA2312 flash at I/O Address 2400
Verifying QLA2312 flash at I/O Address 2400
Flash at I/O Address 2400 Verified Successfully
New version of the QLA2312 flash BIOS at I/O Address 2600 is 1.43
Programming QLA2312 flash at I/O Address 2400


Error erasing flash at I/O Address 2600
```

Figure 39 QLogic HBA Flash

NOTE: During the installation of the BIOS, the error message shown above will be displayed. This message is normal and does not indicate an issue.

7. Remove the disk and reboot the PC. After POST, you should see the following with your new BIOS version:

```
QLogic Corporation
QLA2200 PCI Fibre Channel Rom BIOS Version 1.43
Copyright (C) QLogic Corporation 1993-2001. All rights reserved.
www.qlogic.com
Press <CTRL-Q> for Fast!UTIL
```

NOTE: See the IBM eServer BladeCenter Fibre Channel Expansion Card Installation and User's Guide that can be found in the Docs folder of this Utility CD for additional information. For more information on the IBMUTIL program, either type "ibmutil /?" at the command prompt.

If you need to restore the defaults settings for the HBA, perform the following steps:

a. Press CTRL+Q to enter *Fast!*UTIL when the following BIOS banner displays:
b. Perform the following steps to restore the default settings and reboot the PC. Use the UP ARROW and DOWN ARROW keys to make your selection.
   i. Select **Configuration Settings**, then press ENTER.
   ii. Select **Restore Default Settings**, then press ENTER.
   iii. You are prompted to press any key. Press ENTER.
   iv. Press ESC.
   v. When prompted, press ENTER to save the changes.
   vi. Select **Exit Fast!UTIL**, then press ENTER.
   vii. When prompted to reboot the system, press ENTER.

8 Red Hat Boot Diskette Modification Process

1. Copy the Network Boot Diskette (bootnet.img) and Supplemental Driver Diskette (drvnet.img) files from the installation media (CD or ISO) to a temporary staging area.
```
# cp /mnt/cdrom/images/{bootnet.img,drvnet.img} .
# ls
bootnet.img  drvnet.img
```

2. Create the temporary mount point and two module directories.
```
# mkdir -p mnt/{boot,drv,initrd}
# ls mnt
boot  drv  initrd
# mkdir modules-boot modules-drv
# ls
bootnet.img  drvnet.img  mnt  modules-boot  modules-drv
```

3. Mount the bootnet.img file via the loopback device
```
# mount -o loop bootnet.img mnt/boot
# ls mnt/boot
boot.msg    general.msg  ldlinux.sys  rescue.msg    vmlinuz
expert.msg  initrd.img   param.msg    syslinux.cfg
```

4. Uncompress the initrd.img file to the staging area.
```
# gunzip -c < mnt/boot/initrd.img > tmp.img
# ls
bootnet.img  drvnet.img  mnt  modules-boot  modules-drv  tmp.img
```

5. Mount the uncompressed initrd.img file (tmp.img) via the loopback device.
```
# mount -o loop tmp.img mnt/initrd
# ls mnt/initrd
bin  dev  etc  linuxrc  lost+found  modules  proc
sbin  tmp  var
# ls mnt/initrd/modules
module-info  modules.cgz  modules.dep  pcitable
```

6. Mount the drvnet.img file via the loopback device.
```
# mount -o loop drvnet.img mnt/drv/
```

7. Extract the driver modules from the archive stored in the initrd.img file to the first module directory created in step #2 (modules-boot).
```
# cd modules-boot
# gunzip -c < ../mnt/initrd/modules/modules.cgz | cpio –id
1412 blocks
# ls
2.4.9-e.40BOOT
# ls 2.4.9-e.40BOOT/
3c59x.o     hid.o      lockd.o    scsi_mod.o  usb-ohci.o     vfat.o
eepro100.o  input.o    nfs.o      sunrpc.o    usb-storage.o
fat.o       keybdev.o  pcnet32.o  usbcore.o   usb-uhci.o
# cd ..
```

8. Extract the driver modules from the archive stored in the drvnet.img file to the second module directory created in step #2 (modules-drv).
```
# cd modules-drv
# gunzip -c < ../mnt/drv/modules.cgz | cpio –id
3714 blocks
```

```
# ls
2.4.9-e.40BOOT
# ls 2.4.9-e.40BOOT/{tg3.o,e1000.o}
tg3.o  e1000.o
```

9. Remove device drivers that will not be utilized from the driver modules extracted to the first module directory in step #7. Since the BladeCenter HS20 and BladeCenter HS40 servers do not support 3Com 3c59x, Intel 10/100, AMD PCnet32 or USB-OHCI devices, it is safe to remove the associated device drivers to free much needed space in the image.

```
# cd modules-boot
# cd 2.4.9-e.40BOOT/
# ls 2.4.9-e.40BOOT/
3c59x.o     hid.o      lockd.o    scsi_mod.o  usb-ohci.o     vfat.o
eepro100.o  input.o    nfs.o      sunrpc.o    usb-storage.o
fat.o       keybdev.o  pcnet32.o  usbcore.o   usb-uhci.o
# rm 3c59x.o eepro100.o pcnet32.o usb-ohci.o
# ls
fat.o  input.o    lockd.o  scsi_mod.o  usbcore.o      usb-uhci.o
hid.o  keybdev.o  nfs.o    sunrpc.o    usb-storage.o  vfat.o
```

10. Copy the Broadcom NetXtreme Gigabit Ethernet Adapter and Intel Gigabit Ethernet Adapter drivers from the second module directory (drvnet.img modules) module archive to the first module directory.

```
# cp ../../modules-drv/2.4.9-e.40BOOT/{tg3.o,e1000.o} .
# ls
e1000.o  hid.o    keybdev.o  nfs.o       sunrpc.o  usbcore.o
usb-uhci.o fat.o    input.o  lockd.o    scsi_mod.o  tg3.o     usb-
storage.o  vfat.o
# cd ..
```

11. Recreate the modules archive from the first modules directory (modules-boot).

```
# ls
2.4.9-e.40BOOT
# find . -type f | grep -v modules.cpio | cpio -o -H crc > modules.cpio
1481 blocks
# ls
2.4.9-e.40BOOT  modules.cpio
# gzip -9 modules.cpio
# ls
2.4.9-e.40BOOT  modules.cpio.gz
# mv modules.cpio.gz modules.cgz
```

12. Copy the required driver disk files from the initrd.img image to the modules-boot directory.

```
# cp ../mnt/initrd/modules/{module-info,modules.dep,pcitable} .
# ls
2.4.9-e.40BOOT  modules.cgz module-info modules.dep pcitable
```

13. Remove driver entries from the initrd.img pcitable file that are associated with the device drivers removed in step #9. Remove all lines that make reference to the removed device drivers (3c59x, eepro100, pcnet32, usb-ohci). For example:

```
0x10b7  0x4500  "3c59x"     "3Com Corporation|3c450 Cyclone/unknown"
0x8086  0x1030  "eepro100"  "Intel Corp.|82559 InBusiness 10/100"
0x1022  0x2001  "pcnet32"   "Advanced Micro Devices"

# cat pcitable | grep –v 3c59x | grep –v eepro100 | grep –v pcnet32 >
pcitable.new
```



```
# mv pcitable.new pcitable
```

14. Remove driver entries from the module-info file that are associated with the device drivers removed in step #9. Remove all stanzas that begin with the name of a removed device driver (3c59x, eepro100, pcnet32, usb-ohci). For example:

```
3c59x
        eth
        "3Com 3c590/3c595/3c90x/3cx980"
eepro100
        eth
        "Intel EtherExpress Pro 100B"
pcnet32
        eth
        "AMD PCnet32"
```

15. Remove driver entries from the modules.dep file that are associated with the device drivers removed in step #9. Remove all lines that make reference to the removed device drivers (3c59x, eepro100, pcnet32, usb-ohci).
```
# cat modules.dep| grep –v 3c59x | grep –v eepro100 | grep –v pcnet32 |
grep –v usb-ochi> modules.dep.new
# mv modules.dep.new modules.dep
```

16. Add driver entries present in the drvnet.img pcitable to the initrd.img pcitable file that are associated with the device drivers added in step #10. All entries in the pcitable file must be sorted by PCI ID.
```
# cat ../mnt/drv/pcitable | grep tg3 >> pcitable
# cat ../mnt/drv/pcitable | grep e1000 >> pcitable
```

    At this point use an editor or sort script to sort the pcitable information by PCI ID.

17. Add driver entries present in the drvnet.img module-info file to the module-info file that are associated with the device drivers added in step #10. For blades, we'll be adding the tg3 and e1000 entries which should look something like this:
```
e1000
   eth
   "Intel EtherExpress/1000 gigabit"
tg3
   eth
   "Broadcom Tigon3 ethernet driver"
```

18. Add driver entries present in the drvnet.img modules.dep file to the modules.dep file that are associated with the device drivers added in step #10.

19. Copy the files copied and modified in steps 12-18 to the mounted initrd.img image.
```
# cp module-info pcitable modules.cgz modules.dep ../mnt/initrd/modules
```

20. Unmount the initrd.img from the loopback device.
```
# cd ..
# umount mnt/initrd
```

21. Compress the initrd.img file.
```
# gzip -9 tmp.img
# mv tmp.img.gz initrd.img
```

22. Copy the initrd.img file to the bootnet.img image.
```
# mv initrd.img mnt/boot
```

23. Unmount the bootnet.img.
```
# umount mnt/boot
```

At this point the bootnet.img file should be ready for use in remote deployments.

8.1 Modify or Create the Post-install Driver Diskette Image

1. Create the supplemental driver disk containing drivers that will be used in the post-install environment.

   Up to this point, we have only been concerned with the drivers that will be used during installation (the BOOT environment) to access storage and network devices. Since we are only updating the ServeRAID drivers in this example, this step is easy – we simply make the driver diskette image file (26k4970.img) available to the Red Hat installer by placing it on an accessible storage device (local or remote) and specify the location in the kickstart file (see the Kickstart File Modification section below).

   If we were adding the ServeRAID adapter to an existing driver diskette image, then we would have to follow a process similar to that described in Section I, but instead of adding the device driver and diskette image files to the boot image, we would add them to an already existing driver diskette image file.

2. Kickstart File Modification

   You must specify the location of the supplemental driver diskette image create in step II by placing the following keyword and parameters in the Red Hat kickstart file:

   If we are installing via FTP --
```
driverdisk --source=ftp://path/to/26k4970.img
```
   If we are installing via HTTP --
```
driverdisk --source=http://path/to/26k4970.img
```
   If we are installing via NFS --
```
driverdisk --source=nfs:host:/path/to/26k4970.img
```

   When adding a network device driver (such as bcm5700 and e1000) to the Red Hat installation initial ramdisk image you must perform the following steps:

3. Create the following script (update_modules.sh) and place it in the rh3 directory of the NFS share. Please note that, by default, it is assumed that the driver diskette image you created in section I has the filename ‘dd.img’ and is located in the root directory of the NFS share.
```
#!/bin/sh
# constants
IMAGE_FILE=/dd.img
DEPMOD=/sbin/depmod

# save current directory
CWD=`pwd`
# create temporary directories
mkdir /tmp/mnt-$$
mkdir /tmp/modules-$$
mount -o loop $IMAGE_FILE /tmp/mnt-$$
cd /tmp/modules-$$
if [ ! -e /tmp/mnt-$$/modules.cgz ]; then
    echo "malformed driver diskette image: $IMAGE_FILE"
    exit 1
```



```
fi
gunzip -c < /tmp/mnt-$$/modules.cgz | cpio -id

for moduleDir in *; do
  for moduleFile in `find $moduleDir -name "*.o"`; do
    if [ -d /lib/modules/$moduleDir ]; then
      targetModuleFiles=`find /lib/modules/$moduleDir -name "$moduleFile"`

      if [ ."$targetModuleFiles" = ."" ]; then
        cp $moduleFile /lib/modules/$moduleDir/kernel
      else
        for targetModuleFile in $targetModuleFiles; do
          cp $moduleFile $targetModuleFile
        done
      fi
      # update module dependencies
      $DEPMOD -a $moduleDir
    fi
  done
done

cd $CWD
umount /tmp/mnt-$$

# clean up
rm -rf /tmp/modules-$$
rm -rf /tmp/mnt-$$
```

4. Run “chmod +x update_modules.sh” to mark it executable.

5. Add the following text (with appropriate changes to reflect the names of the files you created) to the %post section of the kickstart file in order to facilitate execution of the script created in step #1.

```
# copy driver diskette image from the root directory of
# the NFS share to the root directory on the target machine
cp /mnt/source/dd.img /mnt/sysimage

# copy module update script from the root directory of
# the NFS share to the root directory on the target machine
cp /mnt/source/update_modules.sh /mnt/sysimage

# Execute the script in the chrooted environment
chroot /mnt/sysimage ./update_modules.sh
```

Appendix A. Firmware Versions

The following tables contain both the minimum versions required deploy on the BladeCenter and its various components, and the current versions at the time this document was published. Please visit the IBM support site for the most current firmware versions for all components.

www-306.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54597

There is also a URL for registering for proactive emails regarding firmware updates.

https://www-1.ibm.com/support/mysupport/us/en/

First time you visit select register. After you have logged into the site, select edit profile and add the servers that you would like receive notification on.

The tables in that appendix will also help translate between build ids and release versions. You can see the current firmware levels from the management module.

All values current as of Feb 28, 2006.

- Management Module

	Minimum Req	Most Current
Version	1.13	*1.19H*
Release Date	2005/02/11	*2005/12/22*
Revision (Main App)	BRET73H	*BRET82H*

Table 2 Management Module BIOS Information

A minimum of BRET82A is highly recommended since it contains several important fixes:

1) Repeating keystrokes fix
2) Gratuitous arp fix
3) Improved failover and firmware transfer handling

- HS20 BIOS – Model 8832

	Minimum Req	Most Current
Version	1.06	1.10
Release Date	2004/08/12	2005/12/13
Build Id	BSE119AUS	BSE124AUS
Build Date	2004/06/11	2005/10/21

- HS20 ISMP – (8832)

	Minimum Req	Most Current
Version	1.07	1.10
Release Date	2004/07/08	2005/12/11
Build Id	BR8T33A	BR8T36A

- HS20 Diagnostics (8832)

	Minimum Req	Most Current
Version	1.03	1.04
Release Date	2004/07/08	2004/11/09
Build Id	BSYT15A	BSYT16A

Table 3 HS20 – Model 8832 Firmware Information

- HS20 BIOS – Model 8843

	Min Req	Most Current
Version	1.03	1.07
Release Date	2005/04/01	2006/01/20
Build Id	BWE115AUS	BWE122AUS
Build Date	2005/03/25	2005/12/08

- HS20 BMC – (8843)

	Minimum Req	Most Current
Version	1.03	1.14
Release Date	2004/12/08	2006/01/20
Build Id	BWBT13A	BWBT23A

- HS20 Diagnostics (8843)

	Minimum Req	Most Current
Version	1.02	1.04
Release Date	2004/12/22	2005/07/18
Build Id	BWYT05A	BWYT08A

Table 4 HS20 – Model 8843 Firmware Information

- HS40 BIOS

	Minimum Req	Most Current*
Version	1.44	1.62
Release Date	2004/08/13	2005/06/16
Build Id	SBJT44AUS	SBJT62AUS
Build Date	2004/08/11	2005/04/22

***Note:** While version 1.44 is the minimum required, version 1.56 fixes the EFI boot order problem referenced in Item 4 in Section 4.5 and is therefore highly recommended.

- HS40 ISMP

	Minimum Req	Most Current
Version	1.24	1.29
Release Date	2004/08/13	2005/06/21
Build Id	BRMK24A	BRMK29A
Build Date	2004/08/11	2005/05/24

Table 5 HS40 BIOS and ISMP Information

- LS20 BIOS

	Minimum Req	Most Current*
Version	1.04	1.05
Release Date	n/a	2005/12/20
Build Id	BKJT18CUS	BKJT18DUS
Build Date	tbd	tbd

- LS20 BMC

	Minimum Req	Most Current
Version	2.01	2.01

Release Date	2005/12/22	2005/12/22
Build Id	BKBT414C	BKBT414C

Table 6 LS20 BIOS and BMC Information

- QLogic HBA Firmware

  See 7 Flash the QLogic HBA BIOS for instructions on how to update the HBA BIOS.

	Minimum Req	Most Current
Version	**1.43**	**1.43**

Table 7 QLogic Firmware Information

Note: we are now recommending v1.43 for the QLogic BIOS. Old minimum was v1.38.

Appendix B. Worksheets

BladeCenter Information

BladeCenter Location	
User ID for BladeCenter Modules	
Password for BladeCenter Modules	

BladeCenter Components

	Firmware Level	IP Address
Management Module 1		
Management Module 2		
Ethernet Switch Module 1		
Ethernet Switch Module 2		
Fibre Channel Switch Module 1		
Fibre Channel Switch Module 2		

Blade Configuration Information

Hostname	IP Address	Location	Operating System

SAN Configuration

System Name	Adapter	WWID	HDC Name	HDC WWID	HDC WWID	HDC WWID	HDC WWID
	1						
	2						
	3						
	4						

Index

